TOOLS TO FACILITATE AUTONOMOUS QUADROTOR-BASED CINEMATOGRAPHY

NIELS JOUBERT
12/12/2016

Thank you Pat, that's incredibly kind of you. Thank you very much everyone for coming, it's great to see so many friendly faces here, it really means a lot. If you're wondering why there's cake, it's Professor Alonso's birthday today! Happy Birthday! Clap guys! He's on my committee, gotta make him feel good! Today I'd like to tell you about an exciting new direction of research in Computer Graphics we've started here at Stanford University. Cinematography is the art, and skill, of moviemaking. I'm going to tell you about the tools we've built to make this art and skill more accessible to people using quadrotor cameras. Let's get into it.

====

For the first time, flying, robotic cameras are ubiquitously available. I'm excited to share a vision of where we, as researchers, might go with this new technology, and the work we've done to take steps into making this vision a reality.

QUADROTOR CAMERAS HAVE HUGE POTENTIAL IN MOVIEMAKING

[3DR Solo, 2015]

The last few years we have seen the proliferation of orientable cameras carried aloft by small quadrotor aircraft, like this one. We call these Quadrotor Cameras, or more colloquially, drones. They have huge potential to revolutionize filmmaking for a couple of reasons.

===

Define Cinematography: The Art and Skill of Moviemaking

QUADROTOR CAMERAS LOWER COST OF AERIAL CINEMATOGRAPHY

Helicopter
Rent at >$1000 per hour

Quadrotor
Purchase for ~$1000

[Red Bull, 2014]

[Cut Media, 2014]

First off, they're inexpensive. Quadrotor Cameras drastically lower the cost of aerial cinematography.

On the left is a mountain biker being filmed from a helicopter. This costs more than $1000 per hour. On the right, you see a mountain biker being filmed from a drone. This type of drone can be *purchased* for around $1000, and capture similarly impressive video.

QUADROTOR CAMERAS SUBSUME OTHER CAMERA RIGS

Furthermore, quadrotor cameras can achieve similar shots to several other camera rigs normally reserved for professional filmmakers.

On the left is a camera crane. Cranes are often used to vertically lift the camera, revealing a scene. On the right is a cable-mounted camera. These are used to move a camera along a predetermined straight line.

Quadrotor cameras can achieve both of these effects and many more.

QUADROTOR CAMERAS ENABLE NEW CINEMATOGRAPHY

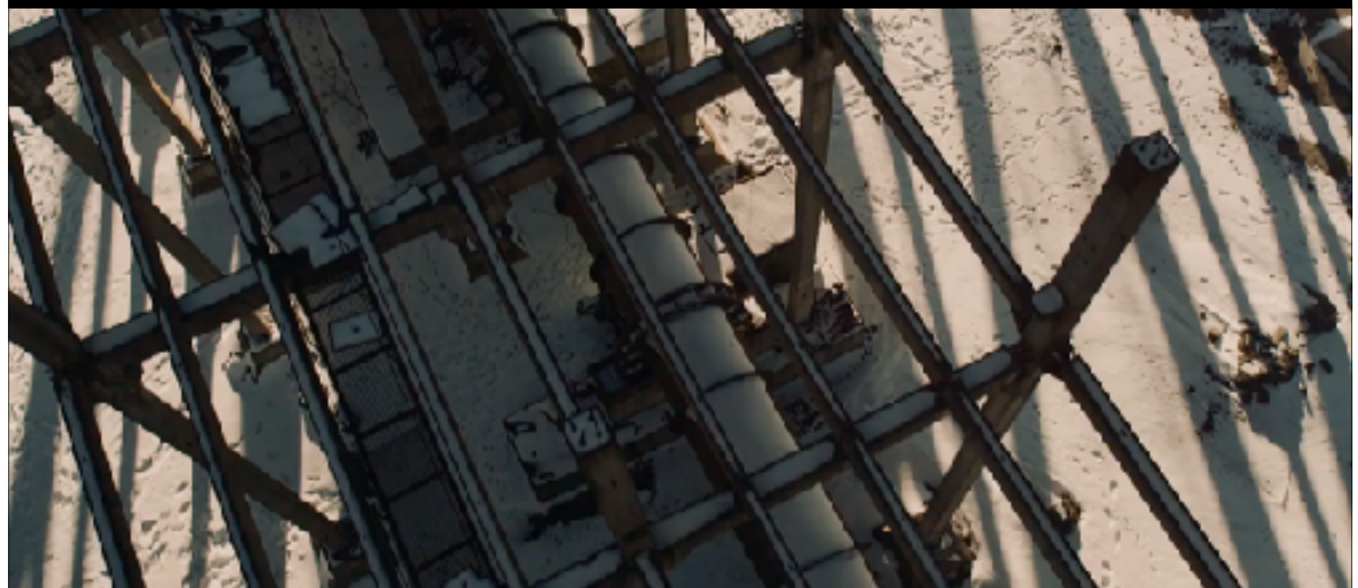

AeroCine, “The Fallen”, 2015
NYC Drone Film Festival Award Winner

And perhaps most exciting, we can capture never-before-seen impressive cinematography with this new camera technology. This is because they can fly to unique vantage points and execute dynamic camera moves in 3D space.

<CLICK>

Here’s a shot from the winning entry in the first-ever Drone Film Festival, held in New York last year. The camera smoothly tilts and yaws while it traces out a complex path through its environment. I’m hard pressed to think of another tool we could use to film this. The only other times I’ve seen shots like this is in 3D animated movies when its all done with special effects

Unfortunately…

BUT REQUIRES CHALLENGING MANUAL OPERATION

[AeroCine]

… it took a crew of two highly-skilled professionals to capture that shot.

On the left is the guy flying the quadrotor, on the right is the guy rotating the camera around.

That was shot by Aerocine. We interviewed them as part of exploratory interviews we conducted. They described how they employ the world’s number one model helicopter pilot to fly their quadrotor, and a camera operator skilled in visual composition to orient the camera. In their words, capturing complex shots is like improvising jazz music, where every member of the group has to be in sync, working together, and anticipating each other’s movements.

This is true for amateurs as well. When you are I want to capture video with a quadrotor, we have to do the job of both the pilot and the cinematographer. And we have to control these cameras manually, in real time…

===

On the right is the controller of a popular consumer quadrotor. It has two main joysticks and a control wheel to control the position and orientation of the camera. I think most everyone can get behind the idea that, using this controller, capturing anything but the most basic shots require a lot of dexterity and skill

Jeff and Oscar

USING INTERFACES NOT DESIGNED FOR CINEMATOGRAPHY

DJI Controller, Peter Homer, 2015

Niels' Grandpa's Controller, 1989

… using a set of joysticks and control wheels.

Here is the controller for the DJI Phantom, the most popular consumer quad rotor camera currently on the market. These interfaces are expressive enough to capture impressive cinematography, as you've seen. But i'd argue that the tools we're using have not been designed with cinematography in mind. They don't make it ease.
<CLICK>
In fact, here's the controller my grandpa bought in 1989, when moving a control stick directly corresponded to moving a mechanical control surface on his model aircraft. Today's controllers still mimic the same fundamental design.

PROBLEM

QUADROTOR CINEMATOGRAPHY
REQUIRES

TECHNICAL SKILL OF FLYING, AND
ARTISTIC SKILL OF COMPOSITION,
SIMULTANEOUSLY APPLIED IN REAL TIME

Let me state this problem more precisely.
<CLICK>
Quadrotor cinematography requires
<CLICK>
the technical skill of piloting,
<CLICK>
and the artistic skill of Visual composition,
<CLICK>
applied simultaneously in real time

The work I’m going to show you is an attempt at improving on this state of affairs.

====

The way I like to phrase this is that we expect the user to work on the level of a camera operator.
We’ll start by asking the question, how can we have people NOT fly these by hand?

APPROACH

CINEMATOGRAPHY-FIRST INTERACTION

We will **automate the role of the pilot**

Provide tools that enables the user to
focus on the composition of shots

And I call this approach “cinematography-first interaction”

In my thesis…
<CLICK>
We’re going to get rid of these clunky controllers, and we’re going to automate flying the quadrotor camera.

<CLICK>
We’ll do this by providing I’m going to building tools that lets us express cinematography naturally.

A lot of the magic lies in exactly *how* we ask users to express their cinematic intent. This is quite tricky, since we have to consider the requirements of cinematographers, and match this to the abilities of quadrotor cameras.

=====

I think of it like, we’re going to take over the complicated parts of flying the plane, and we’re going to do that by giving users cinematic input into the plane, and we do the translation to what that would mean algorithmically to fly the quad.

in other worlds

How do we give novice users superpowers?

OVERVIEW

Compose shots using classic 3D Animation primitives, adapted to respect quadrotor camera physics [SIGASIA 2015]

Horus

A Tool for Shot Planning

Compose shots in real time using visual composition principles from filmmaking [arXiv 2016]

Drone Cinematographer

A Tool for Filming People

Today I'll tell you about two tools we've built…

<CLICK>

First, I'll show you how we can compose shots for quadrotor cameras using classic 3D animation primitives. We have to do just one tricky thing, and that is to adapt these primitives to respect quadrotor camera physics.

We'll use this insights…
<CLICK>
…to build a shot planning tool, called Horus. With Horus, we free you from challenging manual quadrotor control, and the complex real-time decision making it requires. Users cant focus on shot planning and cinematography offline, enabling a new level of fine-tuning and control of quadrotor camera shots.

We don't have to stop here.

<CLICK>

Then I'll shot you how we can compose shots for quadrotor cameras using visual composition principles from filmmaking. In the process, we'll invent a new path planning algorithm, and validate a new person-tracking technology. Using these insights…

<CLICK>

I'll show you how we build the Drone Cinematographer. Using the Drone Cinematographer, you can film people in real time, again free from manual quadrotor control. This tool offers guidance to novice cinematographers in the form of pre-composed shot selection

OVERVIEW

Compose shots using classic 3D Animation primitives, adapted to respect quadrotor camera physics [SIGASIA 2015]

Horus

A Tool for Shot Planning

Compose shots in real time using visual composition principles from filmmaking [arXiv 2016]

Drone Cinematographer

A Tool for Filming People

Let me start off by showing you our solution to the problem of having to manually fly quadrotors.
This is work I did in collaboration with Mike, Anh, Floraine and Pat.

Before we get into any heavy math, let me show you what we can do with this tool…

SHOT WITH HORUS

This shot was captured fully autonomously using Horus. It was created by a professional cinematographer from Weta Digital, the studio behind Lord of the Rings

MISSION PLANNING TOOLS ALSO AUTOMATE THE PILOT

Mission Planner, ArduPilot, 2015

We're not the first to consider getting rid of those controllers and automating the pilot. The robotics community have developed mission planning tools, such as the one I'm showing you here. You can program a quadrotor to follow a path and do all kinds of actions by programming in waypoints on this map, and creating this list of commands you see at the bottom…

BUT THEY DO NOT ALLOW USERS TO FOCUS ON SHOT COMPOSITION

Mission Planner, ArduPilot, 2015

But i think everyone can get behind the idea that these tools are also not designed with cinematography in mind, and they also don’t help with composing shots.

EXPLORATORY STUDY REVEALS CINEMATOGRAPHY CONCERNS

6 professional quadrotor photographers and videographers

Novice, Intermediate, and Expert skill levels

Are familiar with mission planning tools, but have not adopted them for cinematography.

To better understand the needs of cinematographer, we ran an exploratory study interviewing
[CLICK]
six professional quadrotor photographers and videographers. All participants fly quadrotors manually.

[CLICK] and spanned the range of skill levels,

They've all also tried out mission planner software like the one I just showed you, but none of them have adapted it for cinematography.

From this study we extracted requirements for a shot planning tool.

HOW DO WE ALLOW USERS TO…

plan shots visually

precisely control shot timing

rapidly iterate on shots

see **accurate visual previews**

enable **autonomous capture**

From this study we extracted requirements for a shot planning tool. How do we allow users to do the following:

<CLICK>
Plan the shot by setting up what's visually in the frame of the camera
<CLICK>
Precisely control how the camera speeds up and slows down
<Click>
Easily make changes to their shot
<Click>
Preview what the shot will look like
<CLICK>
Capture it faithfully

USE 3D ANIMATION PRIMITIVES!

plan shots visually

precisely control shot timing

rapidly iterate on shots

see **accurate visual previews**

enable **autonomous capture**

[Real-Time Cameras, Haigh-Huchinson, 2009]
[Christie et al, 2008]

When we create 3D animations in Computer Graphics, we build a 3D environment and then fly a little virtual camera through it using all kinds of specialized camera control algorithms.

What about using primitives from 3D Animation to fly a quadrotor?

<CLICK THROUGH EACH POINT>

USE 3D ANIMATION PRIMITIVES!

- **plan shots visually** ✓
- **precisely control shot timing** ✓
- **rapidly iterate on shots** ✓
- see **accurate visual previews**
- enable **autonomous capture**

[Real-Time Cameras, Haigh-Huchinson, 2009]
[Christie et al, 2008]

These camera planning algorithms allow users to plan shots visually,

precisely control shot timing,

and rapidly iterate on shots…

USE 3D ANIMATION PRIMITIVES!

- **plan shots visually** ✓
- **precisely control shot timing** ✓
- **rapidly iterate on shots** ✓
- see **accurate visual previews** ✗
- enable **autonomous capture** ✗

[Real-Time Cameras, Haigh-Huchinson, 2009]
[Christie et al, 2008]

unfortunately they aren’t built for quadrotors, so they won’t show you accurate visual previews of your shot, and they can’t guarantee that we can actually control a quadrotor in real world to go capture the shot. Let me give you some intuition about why this is the case…

Here, in red, is a camera path your might set up in some 3D animation software.

This seems innocent enough. But let’s take a look at what can go wrong when we try to fly it with a quadrotor.

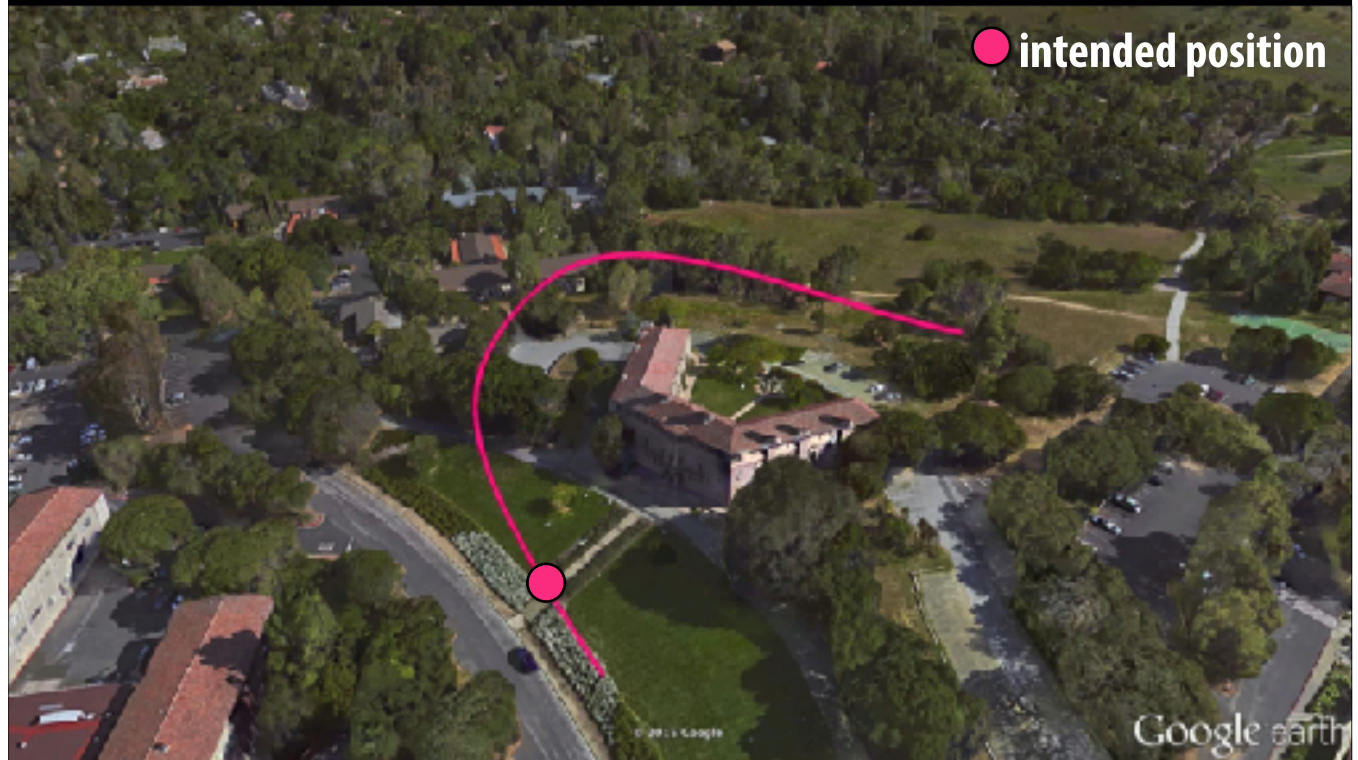

Let’s draw the intended position of our camera using a pink dot.

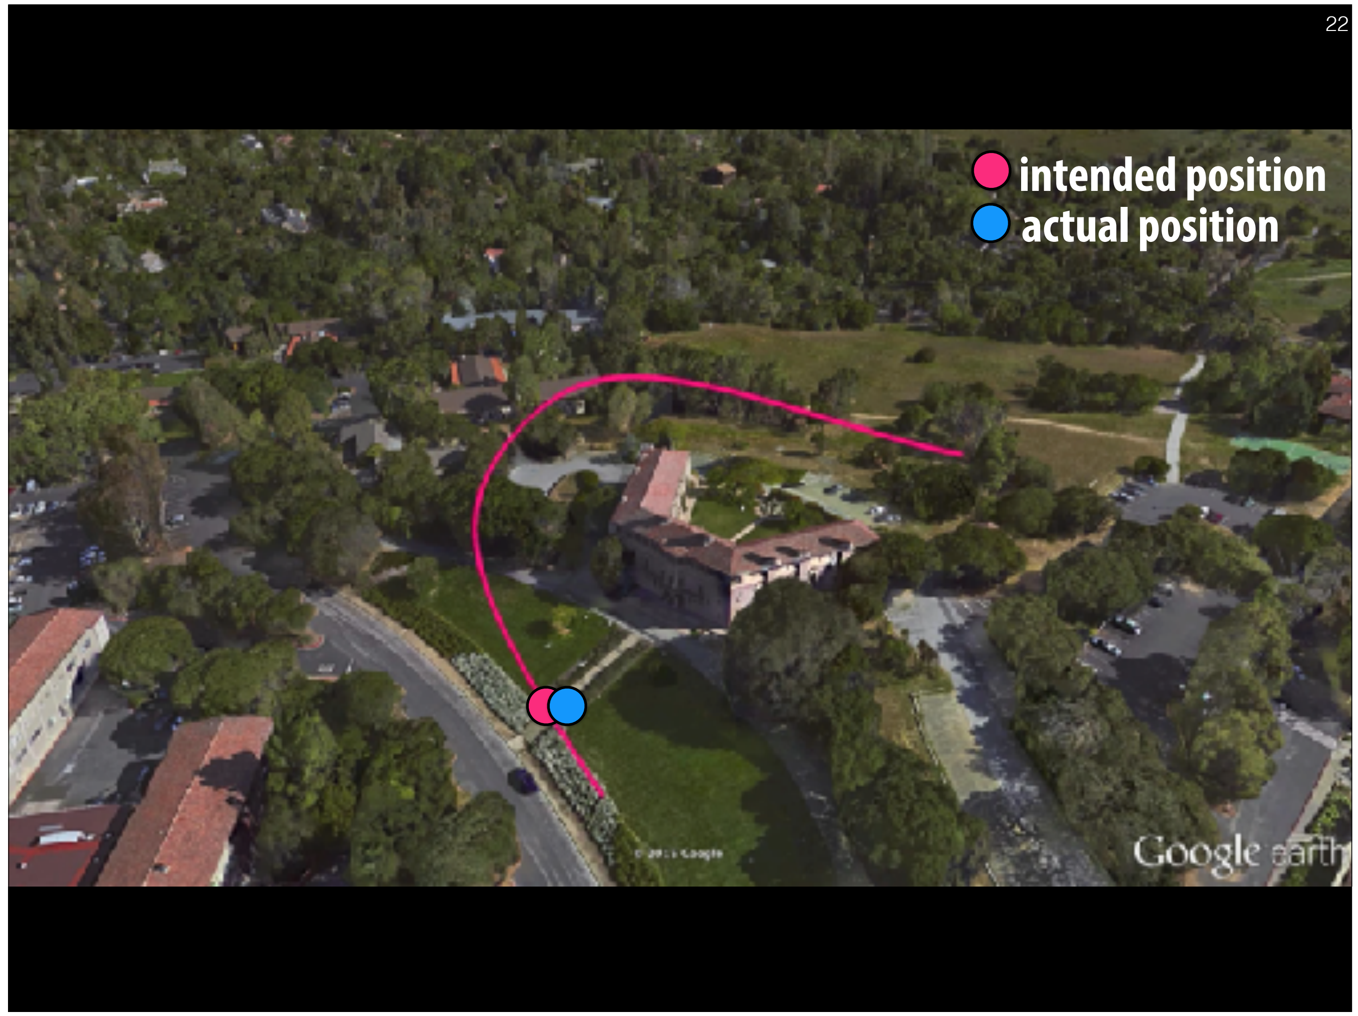

And let’s draw the quadrotor’s actual position with a blue dot.

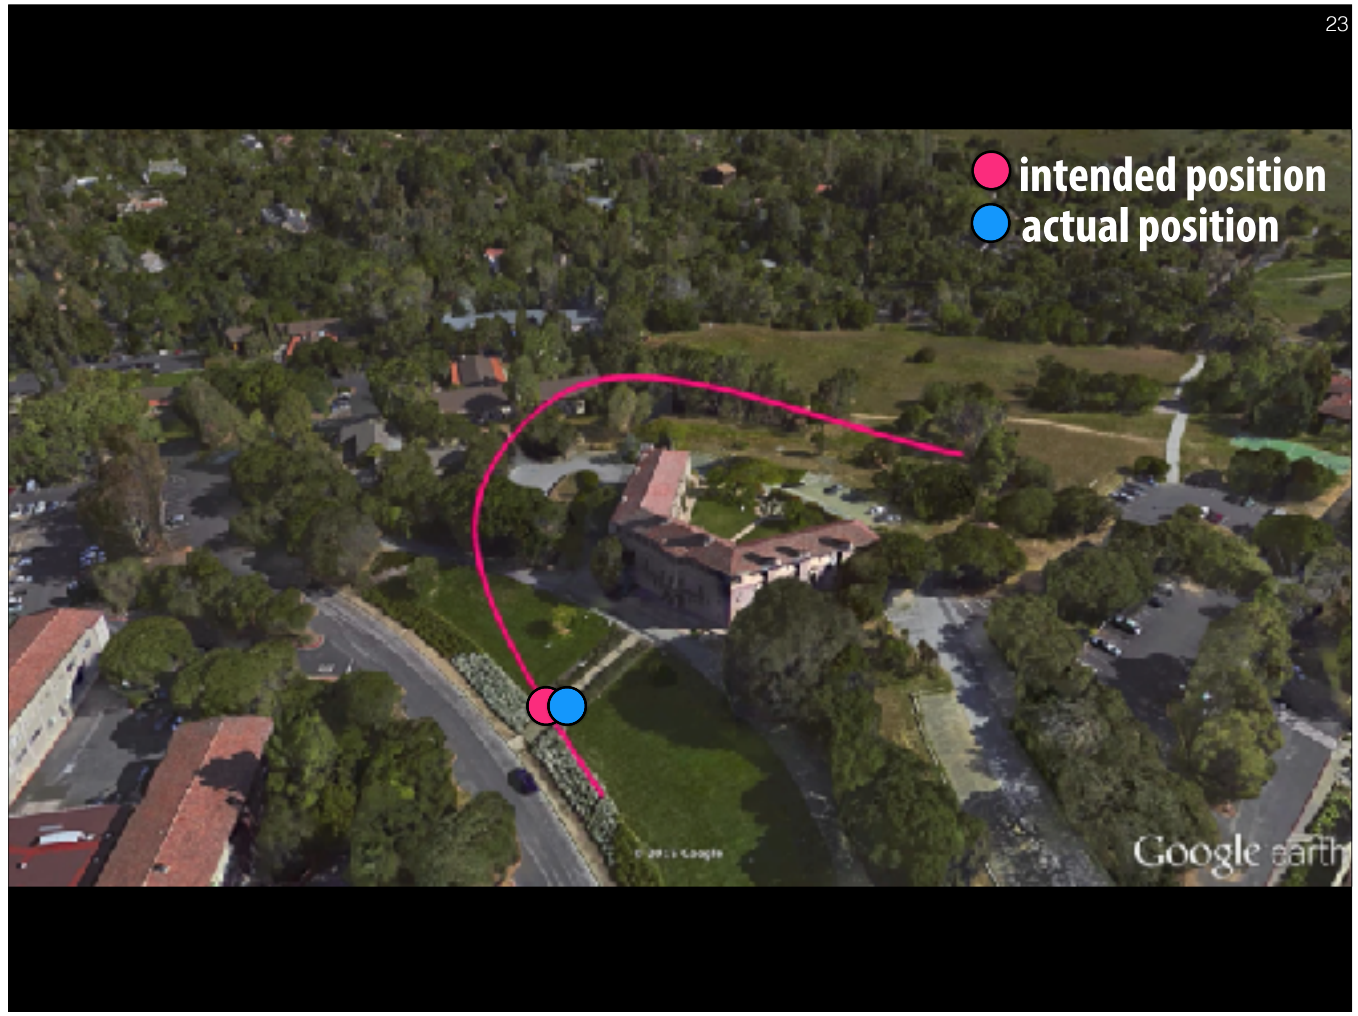

Let’s say the this is a nice and slow, dramatic camera movement. Well that seems to work…

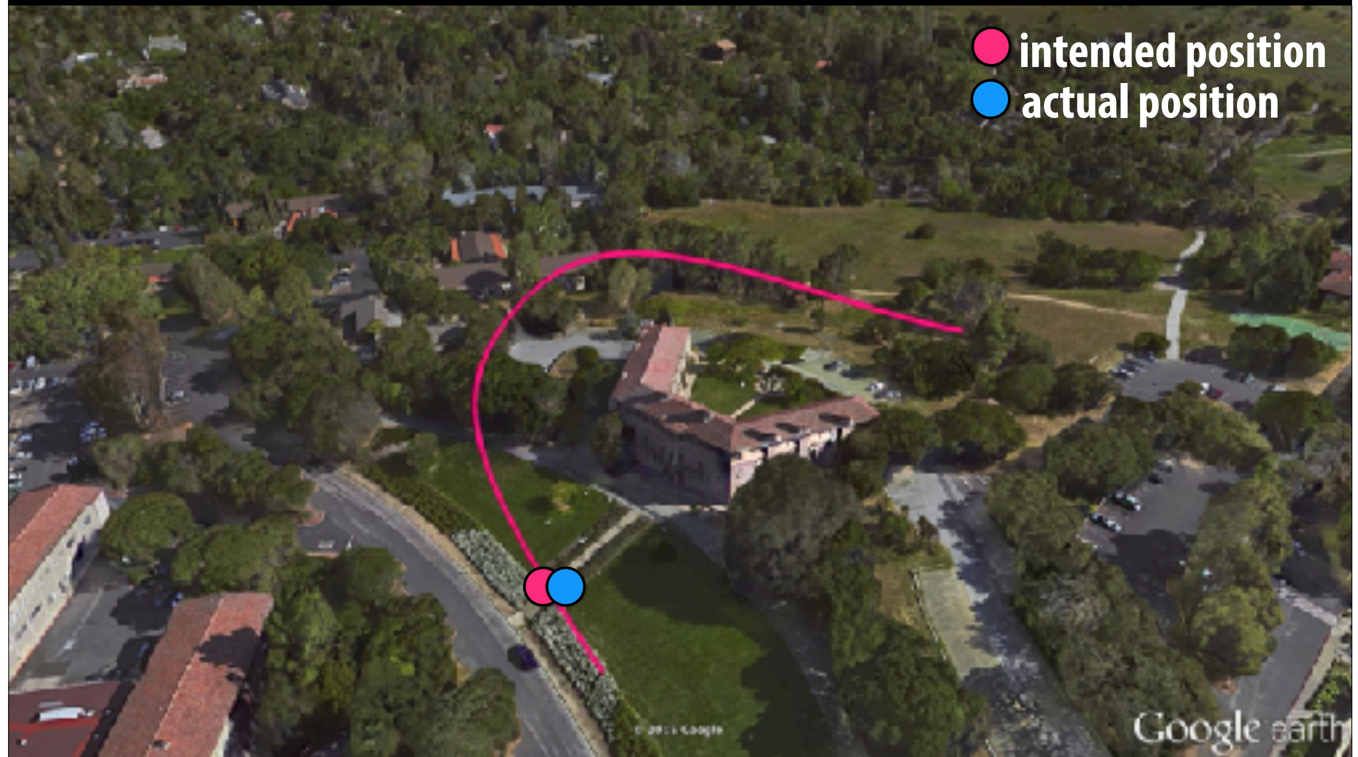

But what if this is a very fast and exciting shot, and the virtual shot we created moves too fast for the quadrotor. Well, your quadrotor can't keep up!

In this case…

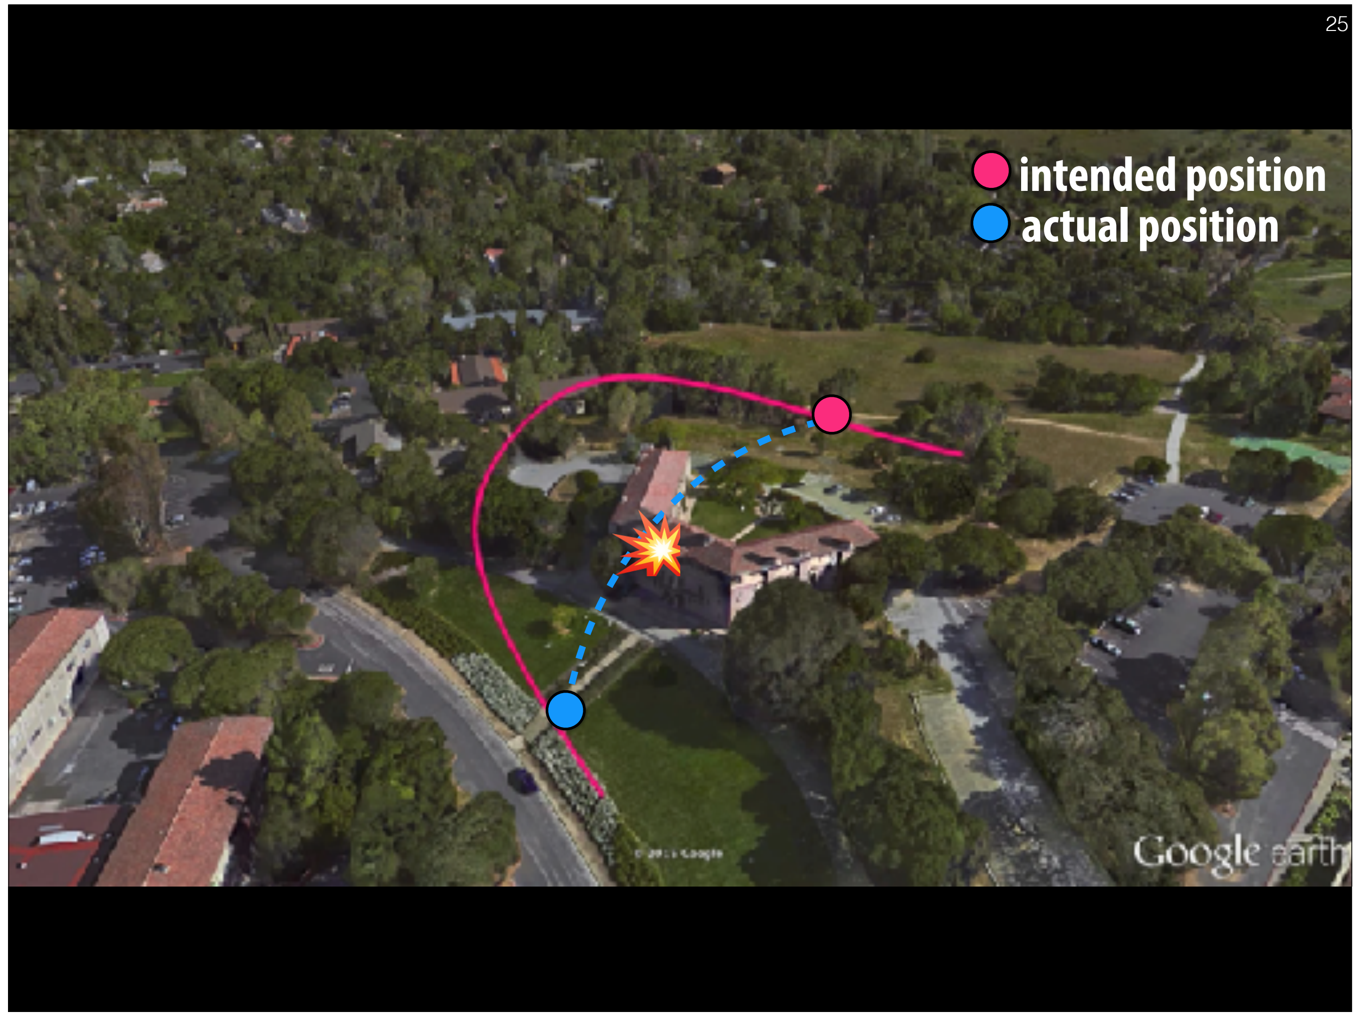

.. the quadrotor will cut corners to try to catch up, and could easily crash into this building, or something else in the environment.

So if we just use a 3D animation tool to create a shot, we won't be able to faithfully capture the shot. Since we thought the quadrotor would follow the fast pink dot, the virtual preview we rendered would be wrong.

So if we don't take the dynamics of the quadrotor into account, our virtual preview might be wrong.

So…

APPROACH

Invent a **shot planning interface**
based on concepts from 3D Animation.

Adapt 3D Animation Primitives to respect quadrotor physics.

[Real-Time Cameras, Haigh-Huchinson, 2009] [Christie et al, 2008]

USE 3D ANIMATION PRIMITIVES!

- **plan shots visually** ✓
- **precisely control shot timing** ✓
- **rapidly iterate on shots** ✓
- see **accurate visual previews** ✓
- enable **autonomous capture** ✓

Now we can satisfy everything, and

USE 3D ANIMATION PRIMITIVES!

- **plan shots visually** ✓
- **precisely control shot timing** ✓
- **rapidly iterate on shots** ✓
- see **accurate visual previews** ✓
- enable **autonomous capture** ✓
- **provide feedback** about quadrotor behavior ✓

we can also provide you with feedback about our quadrotor behavior. This turns out to be quite useful.

THIS IS HORUS

PLAN SHOTS VISUALLY IN A 3D ENVIRONMENT

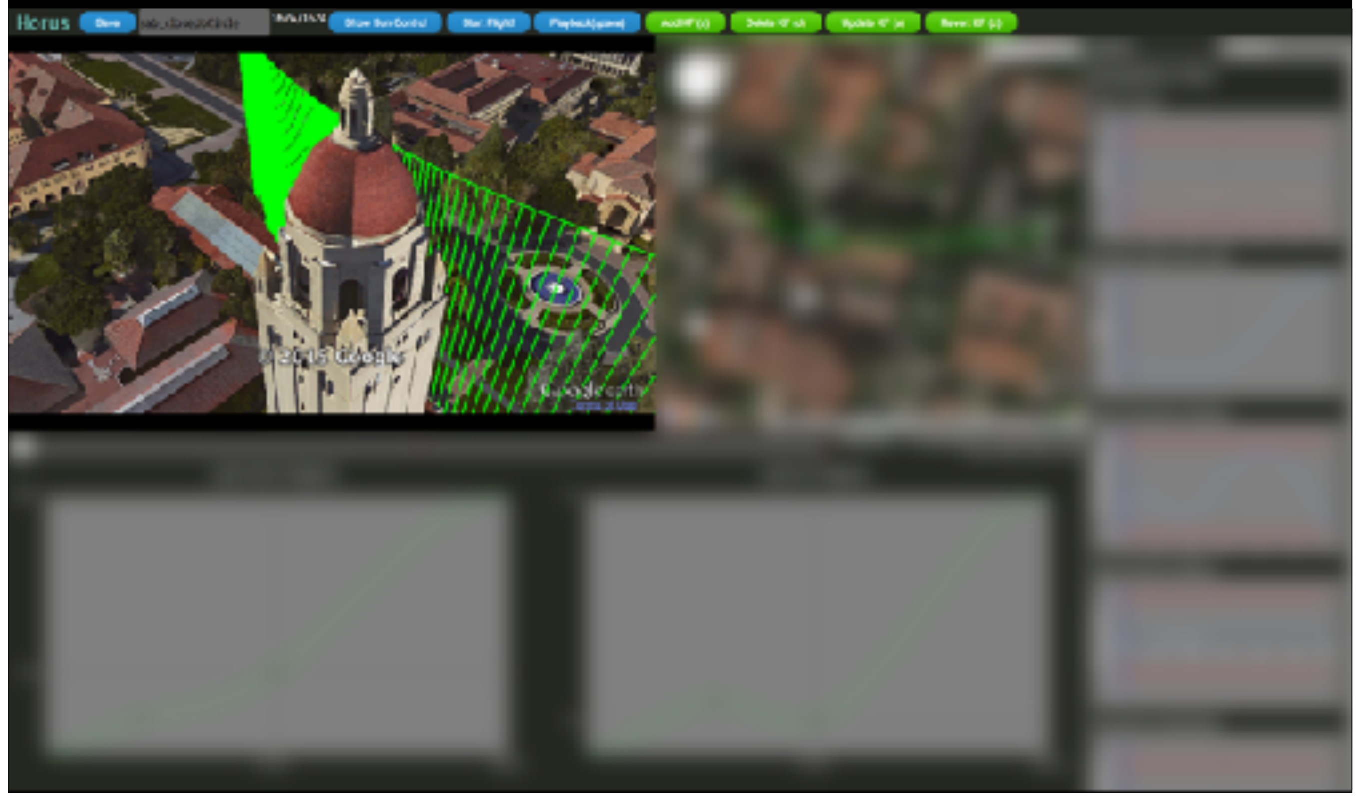

Horus provides a virtual environment for shot planning. We use the Google Earth dataset so we have pretty decent 3D data for the entire US and many parts of the world.

PLAN SHOTS VISUALLY USING SPARSELY SPECIFIED KEYFRAMES

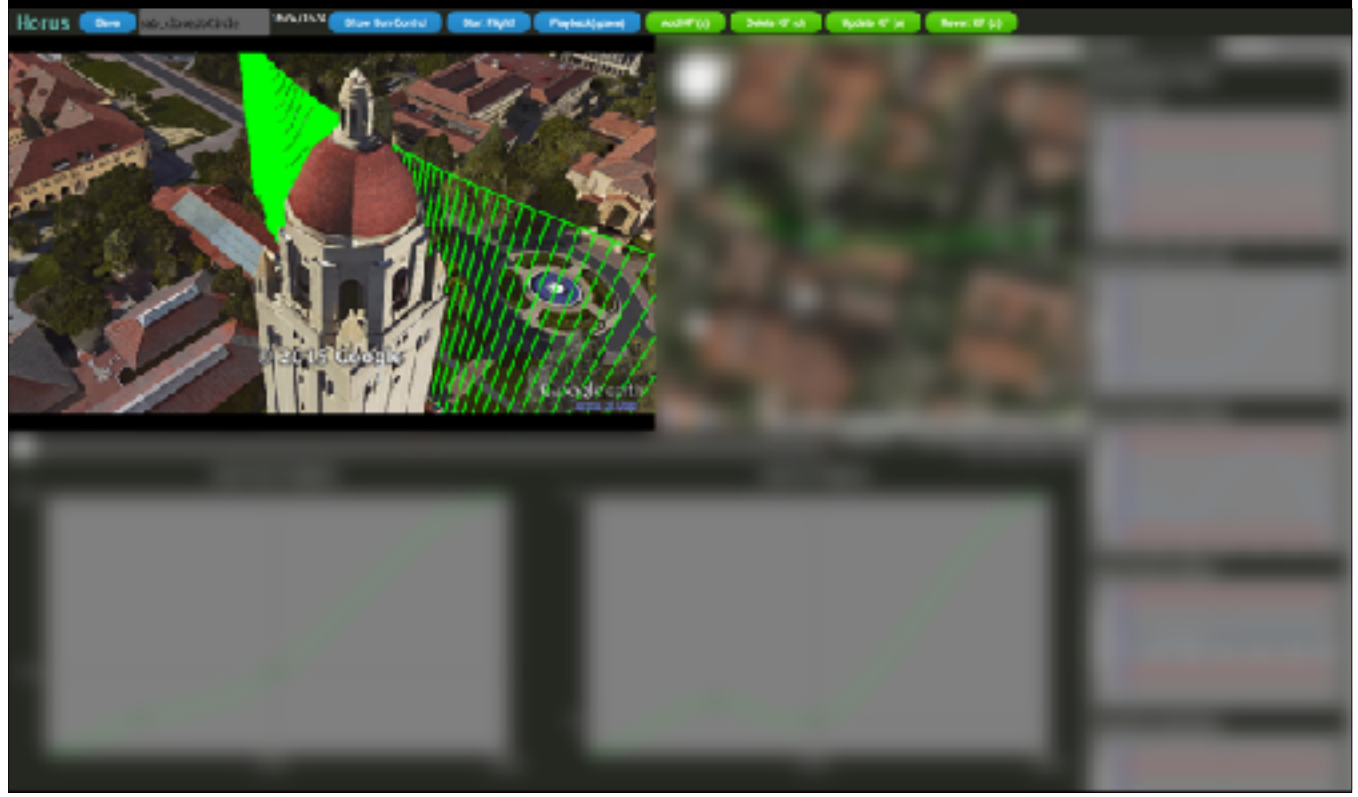

The user can then place a virtual camera at a few representative viewpoints in this environment, and record these viewpoints as keyframes. Each keyframe consists of two points: the look-from point, which is the position of the camera, and the look-at point.

PLAN SHOTS VISUALLY USING SPARSELY SPECIFIED KEYFRAMES

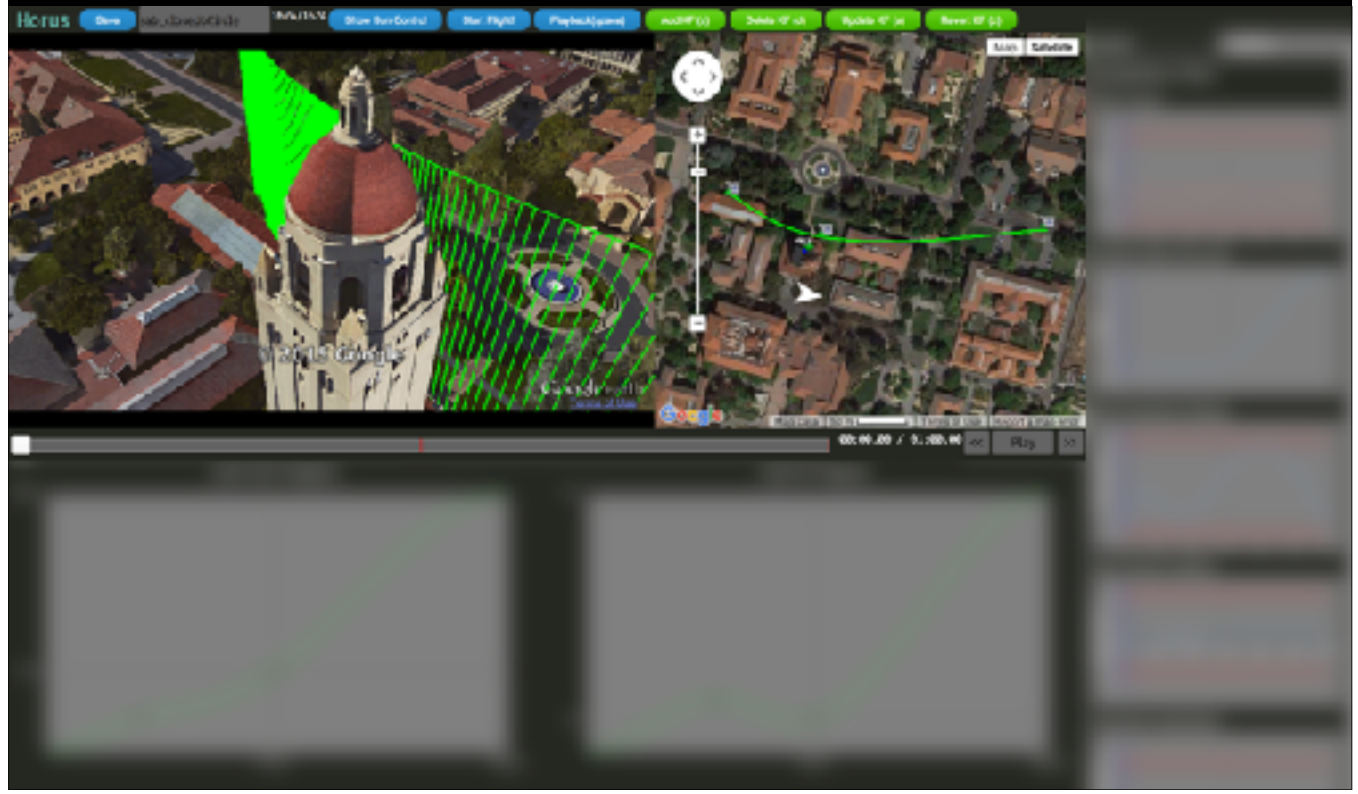

We turn these keyframes into a shot by interpolating a path between their look-from points and look-at points. The real magic lies in how we do this to make sure the quadrotor camera can follow this path. I'll dig into that in a second.

In green you can see the camera trajectory, in a top-down 2D view and in your 3D camera view. One way to fine-tune your shot is to drag these trajectories around.

PRECISELY CONTROL TIMING WITH PROGRESS CURVES

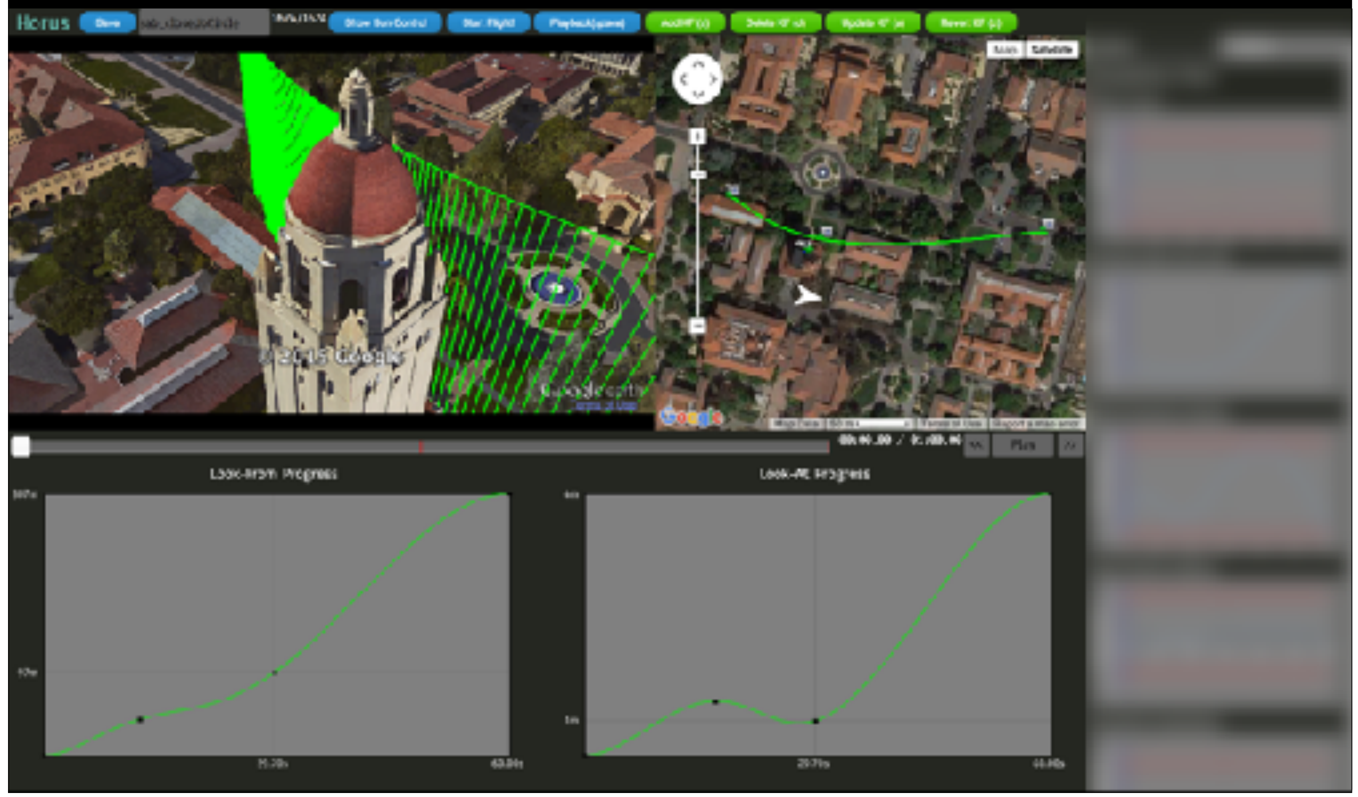

Horus lets the user precisely control the timing of the camera along this path by editing these two progress curves at the bottom. We have a progress curve for the look-from path, and the look-at path. A progress curve encodes how quickly the camera moves along its path as time goes by. So, on the X axis of these curves is time, and on the Y axis is progress along the path.

VISUAL FEEDBACK REGARDING QUADROTOR BEHAVIOR

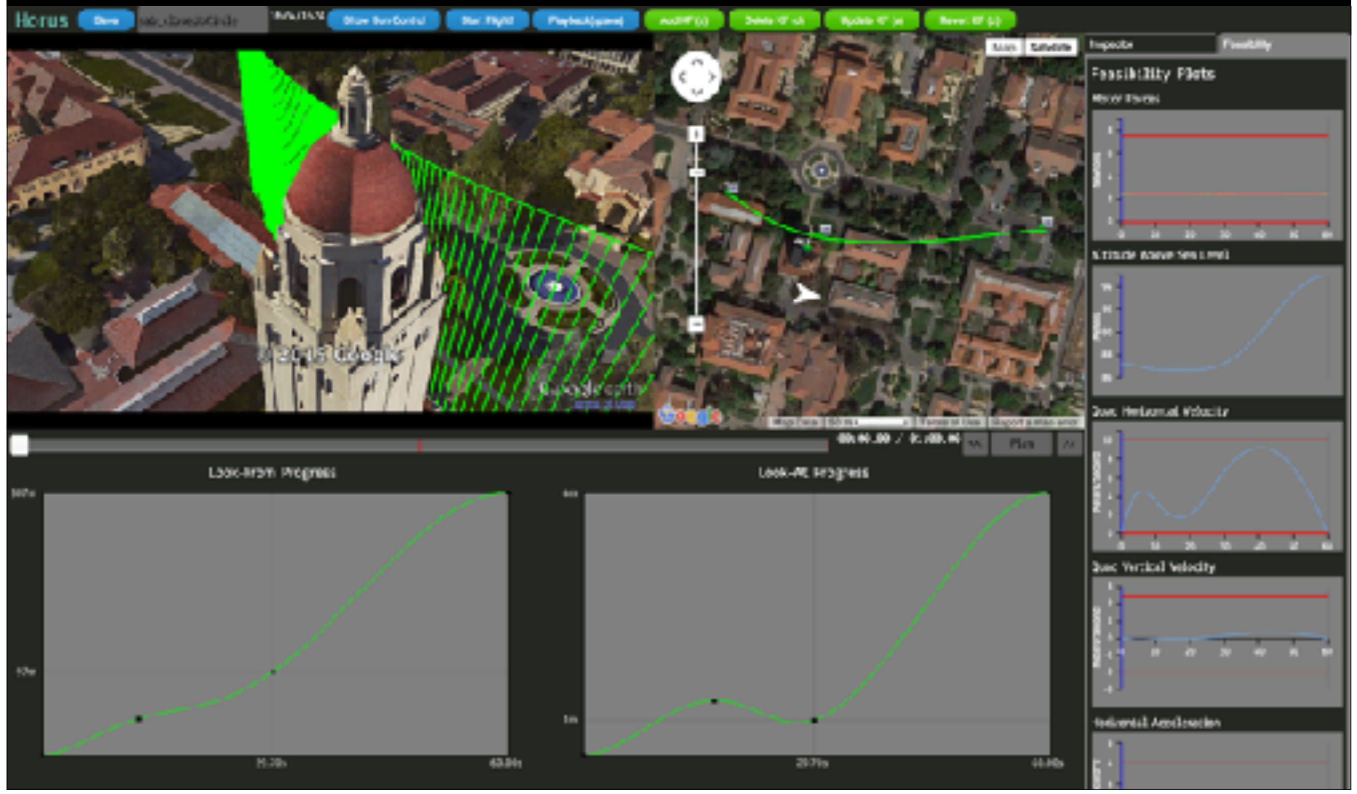

On the right, is a set of feasibility plots. They display the quadrotor's physical behavior along this path, which we get from running a physical simulation in real time.

If your shot were to move too fast, or if they were unattainable in some other way, these plots will show you.

PREVIEW SHOT BEFORE FLIGHT

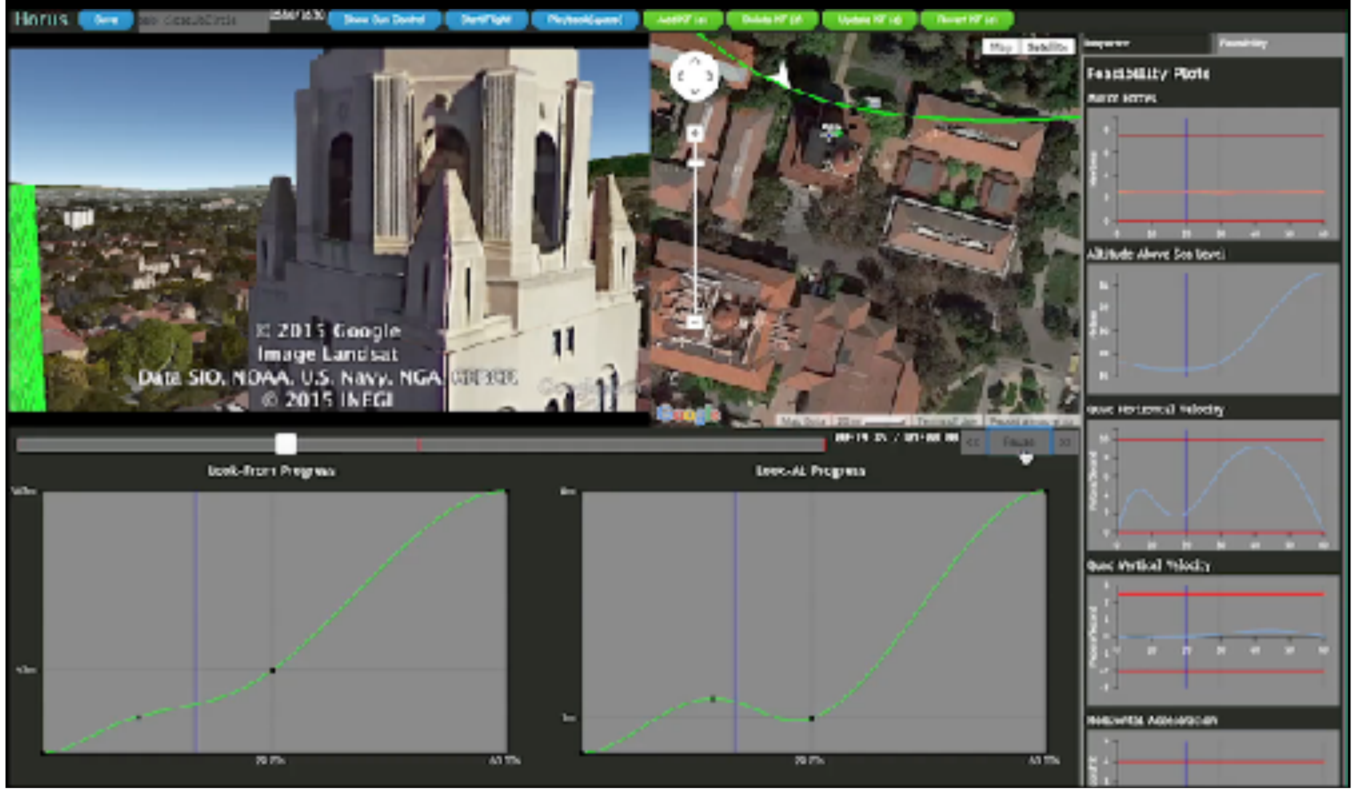

Horus renders a preview of the shot on demand. Using a scrubber, the user can preview their shot and tweak it until they're satisfied. We calibrate this preview to match the camera you're using.

ONE-CLICK AUTONOMOUS CAPTURE

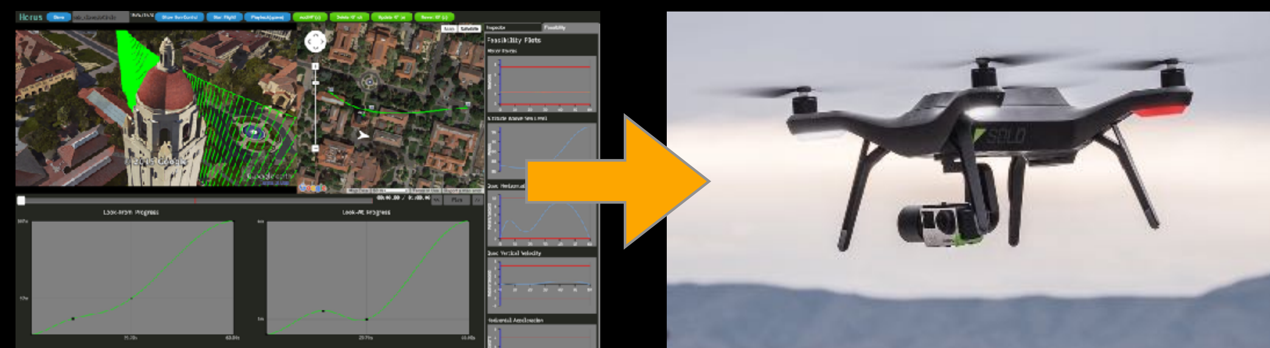

By respecting the physics of quadrotors, Horus can create camera paths that we can capture fully autonomously. All the user has to do is click “go”. During capture, the user can pause, abort, or even tweak the timing of the shot.

TECHNICAL APPROACH

So how do we respect the physics of quadrotor cameras?

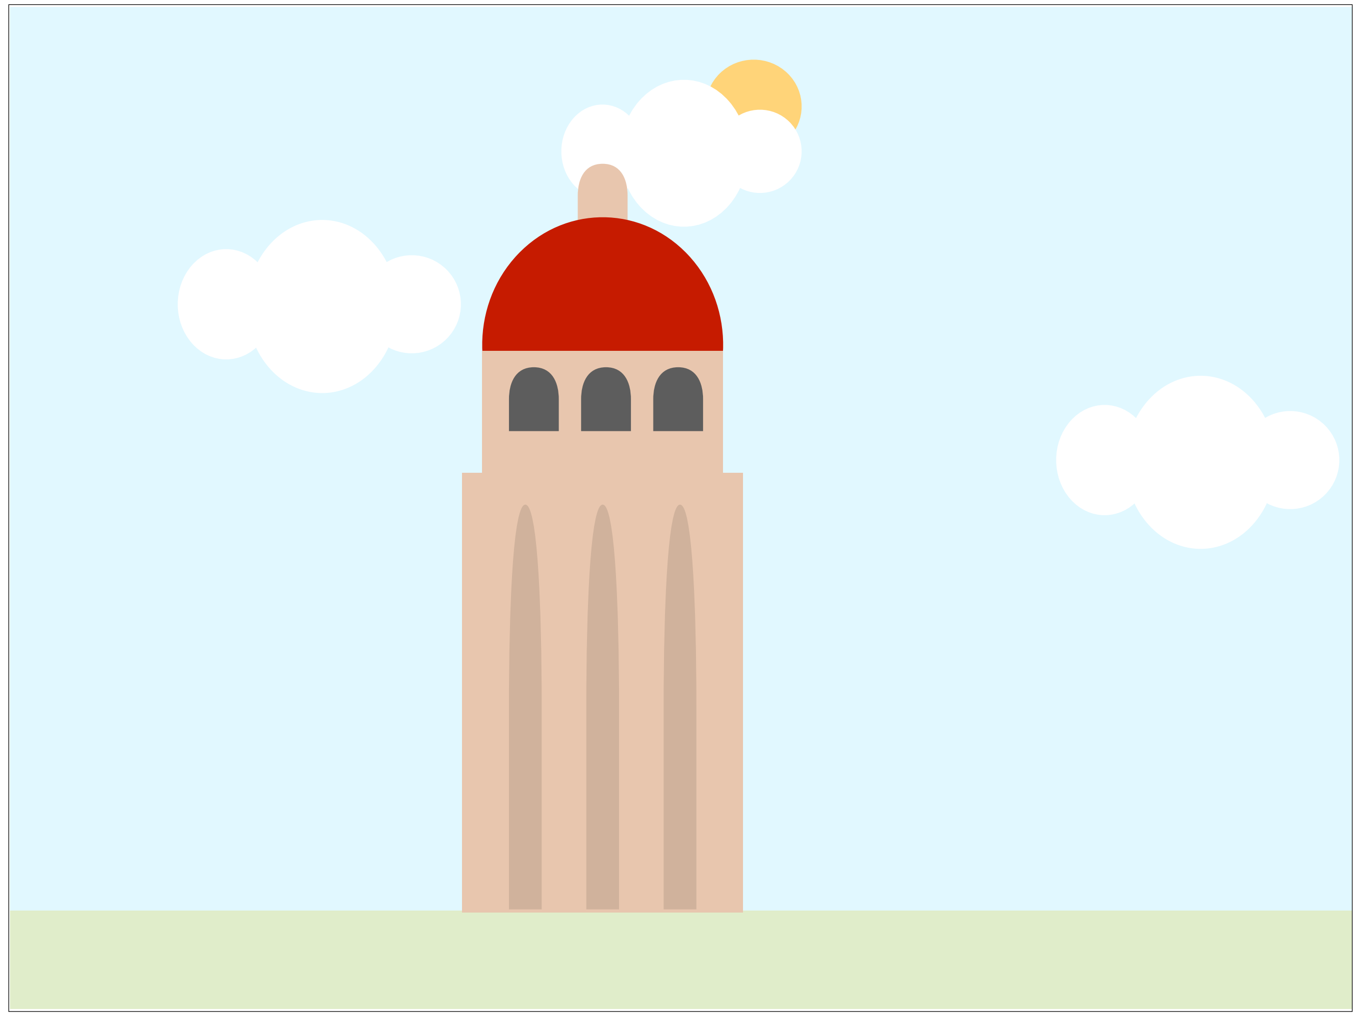

Let me draw you a little cartoon of our problem. It's a cartoon, so I'll oversimplify, but we'll get the important points across without filling up the board with equations. In Horus, the user sets up

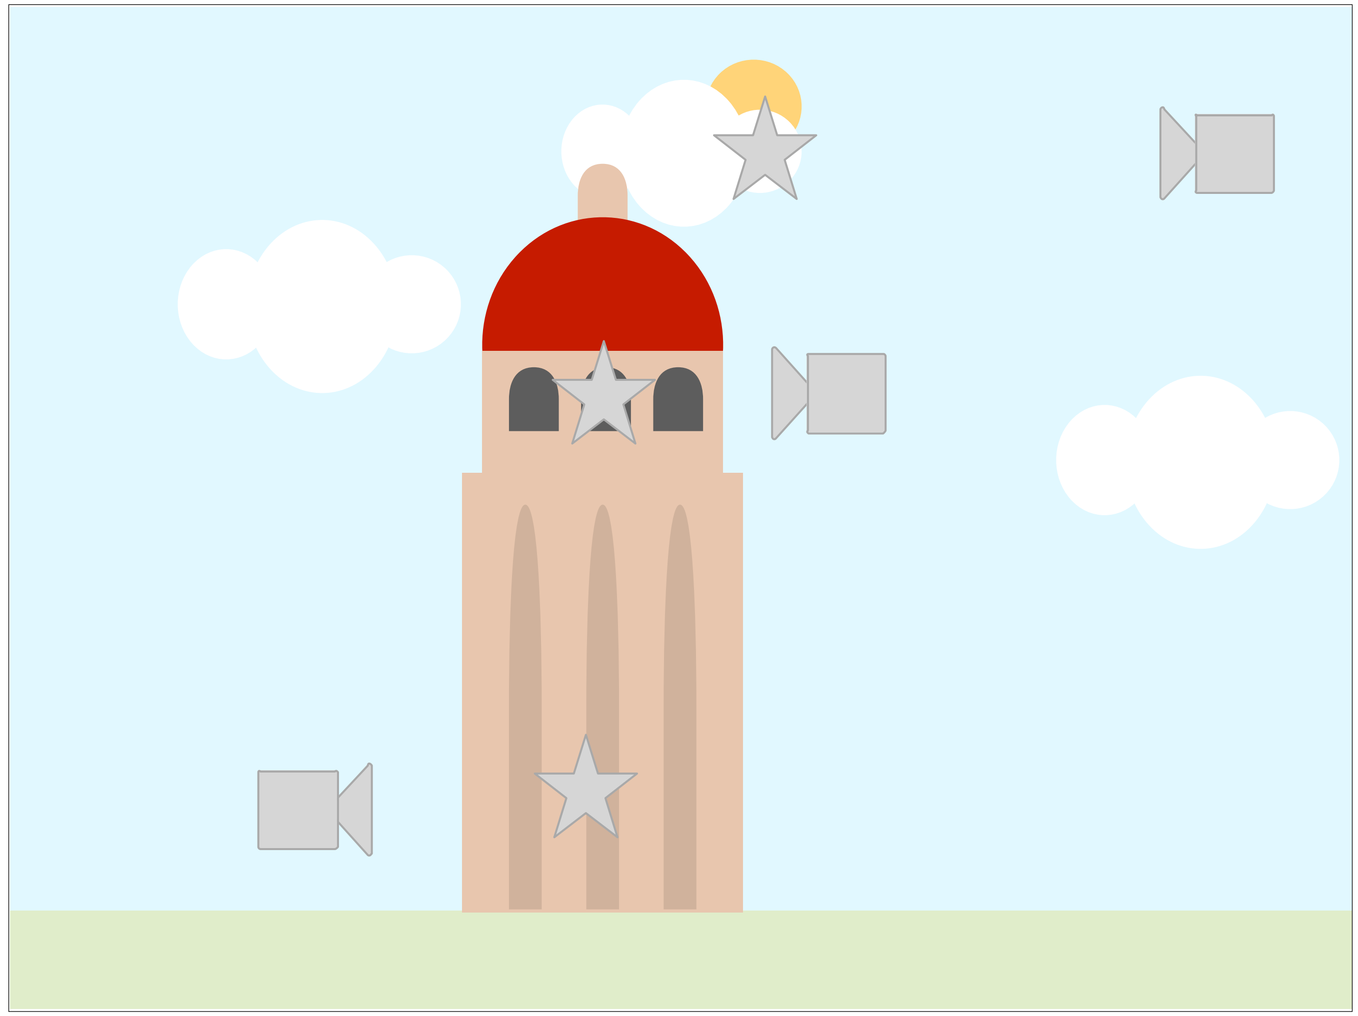

keyframes. Each keyframe consists of a look-from point (the cameras), and a look-at point (the stars)

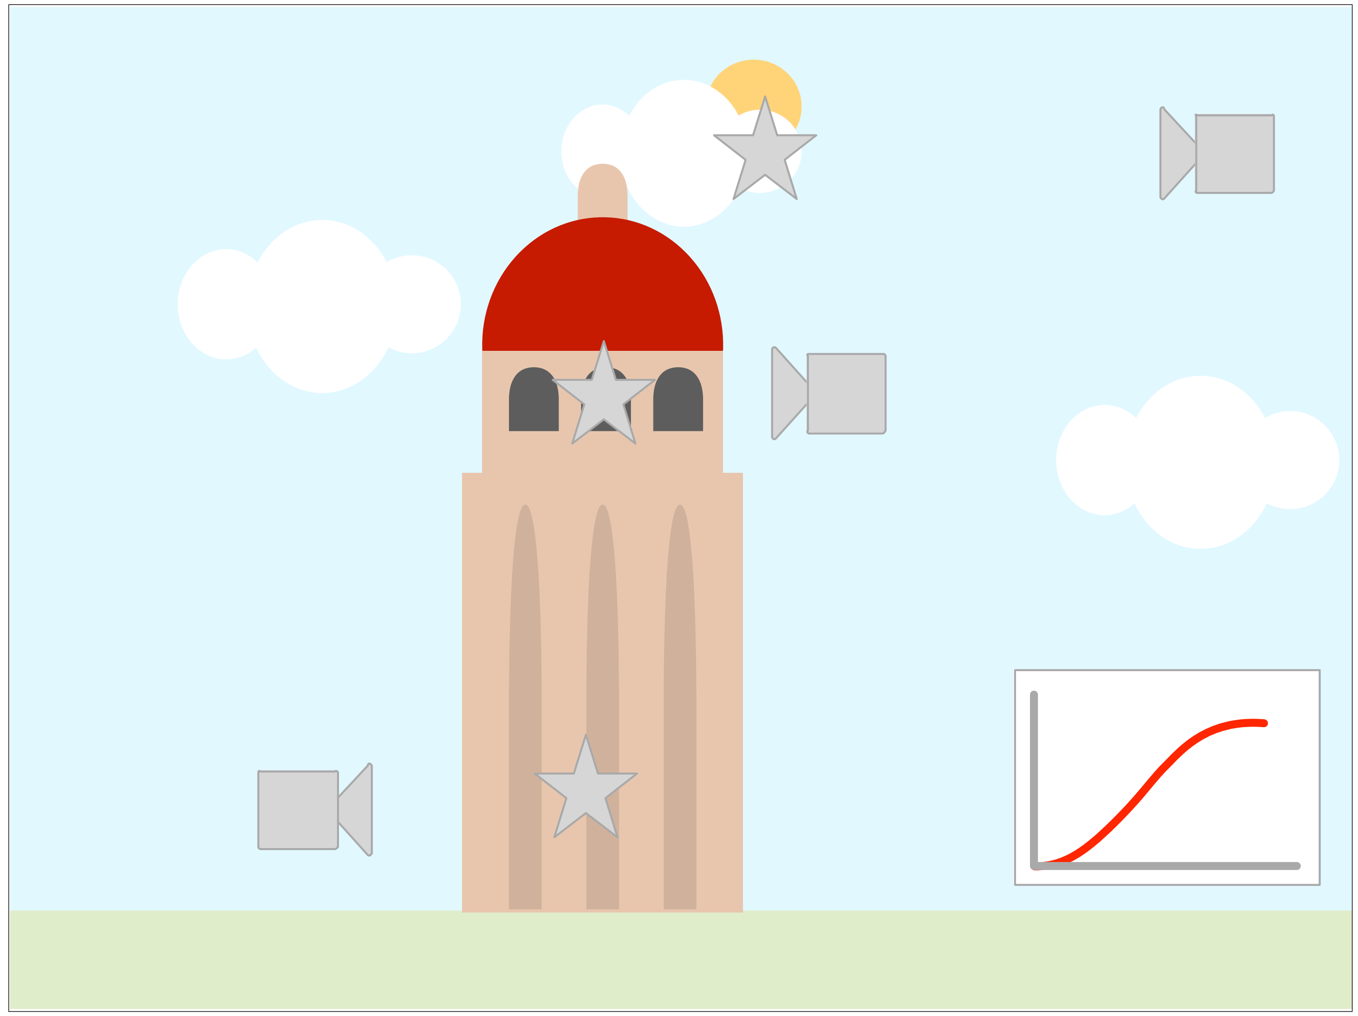

the user also sets up a progress curve. This tells us exactly how quickly we need to move the camera between these points.

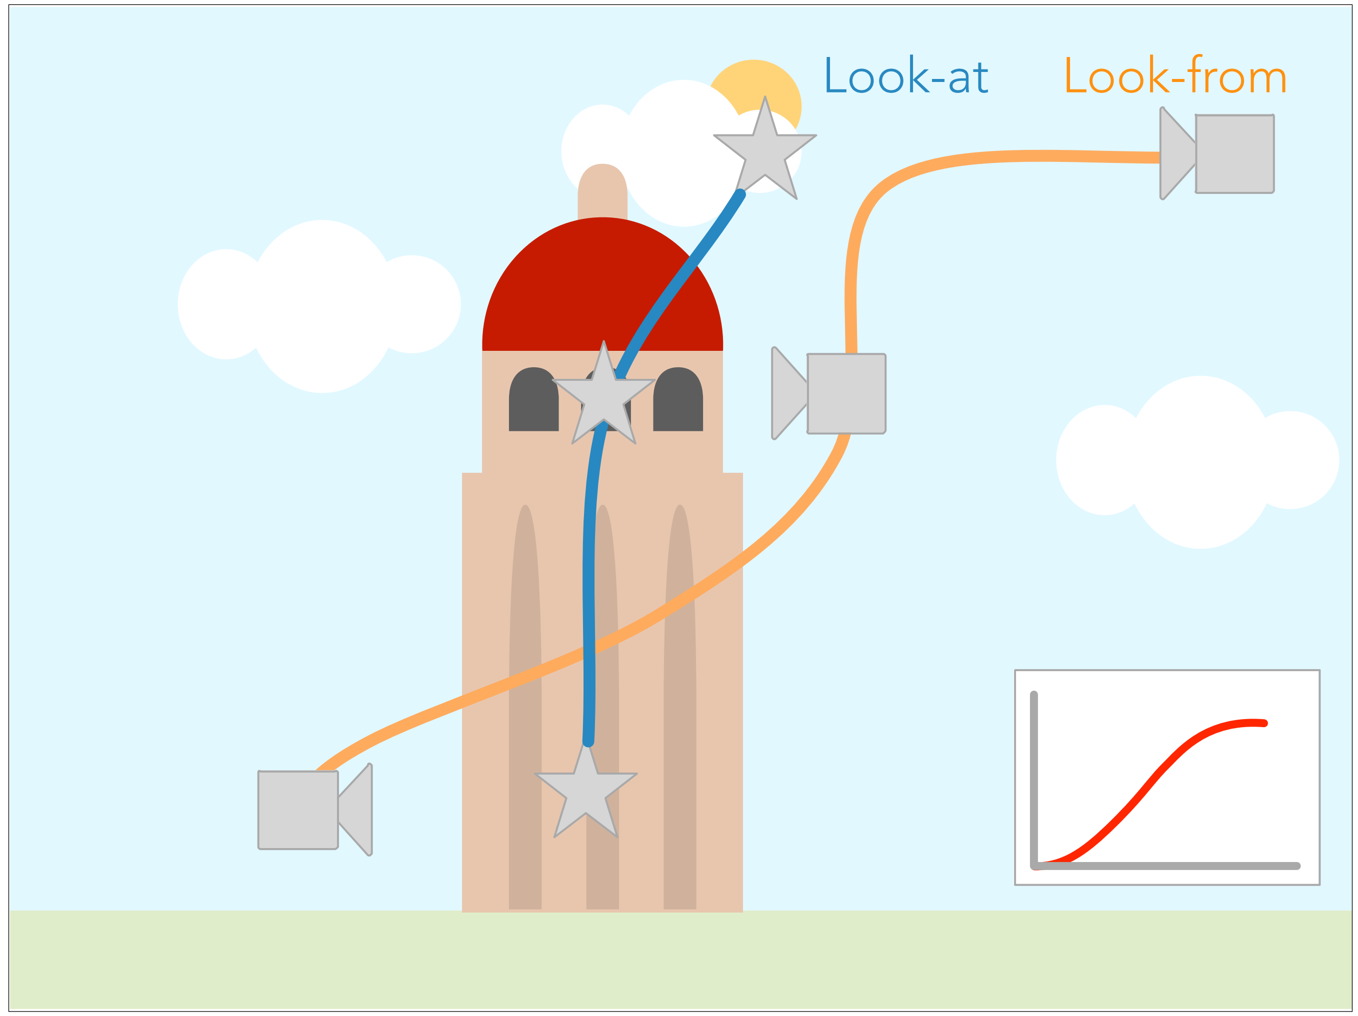

We now want to find a look-from trajectory and a look-at trajectory that interpolates these keyframes, and that moves the quadrotor according to the progress curve.

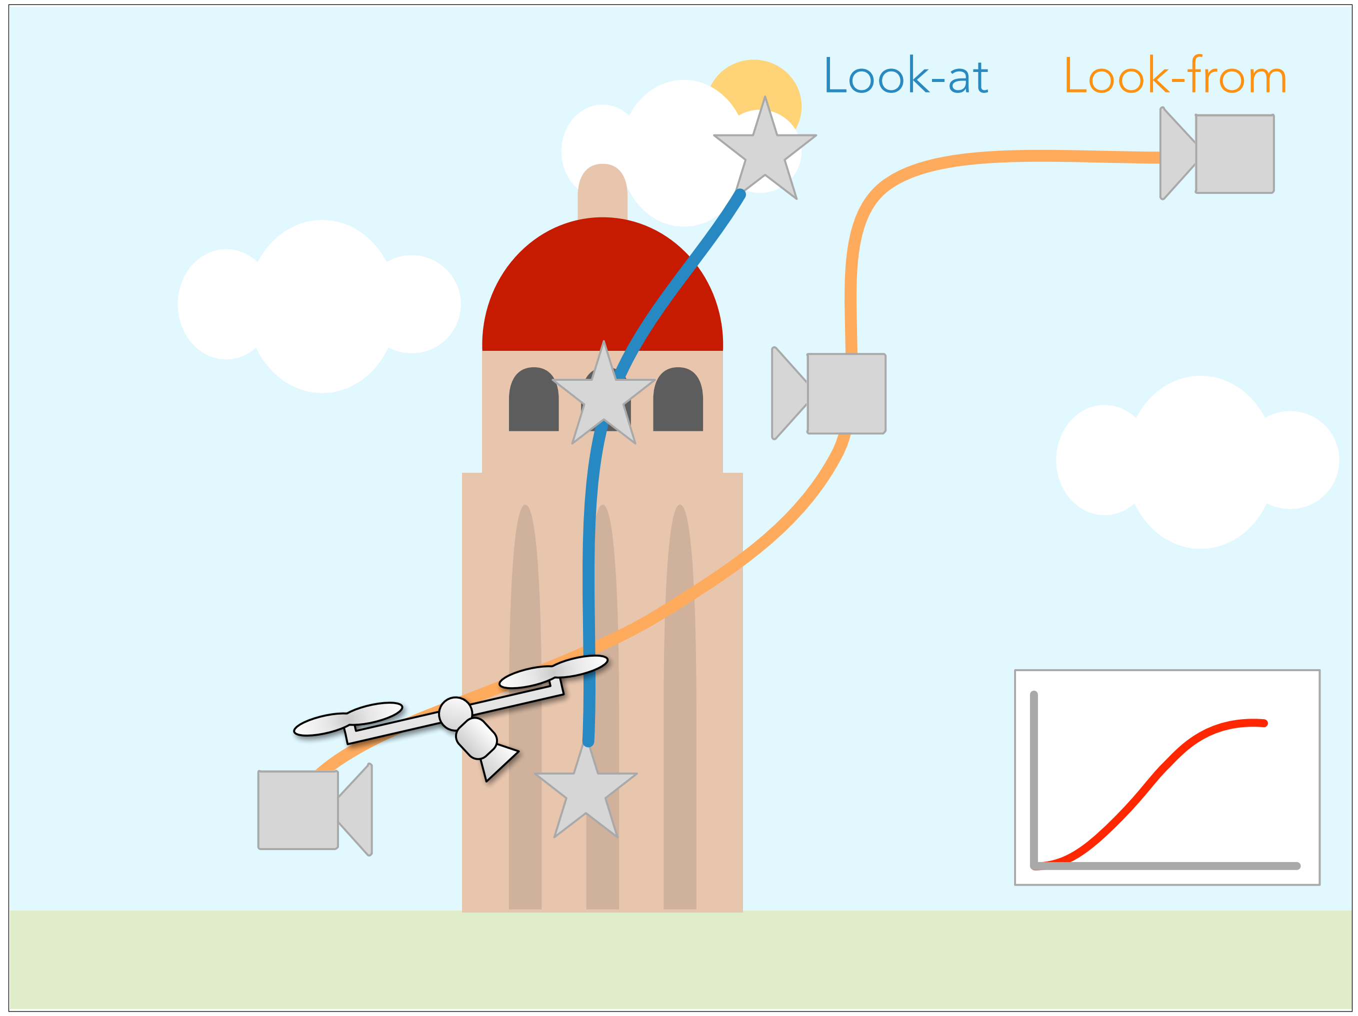

To be able to fly these trajectories, we have to make sure its possible for the quadrotor to …

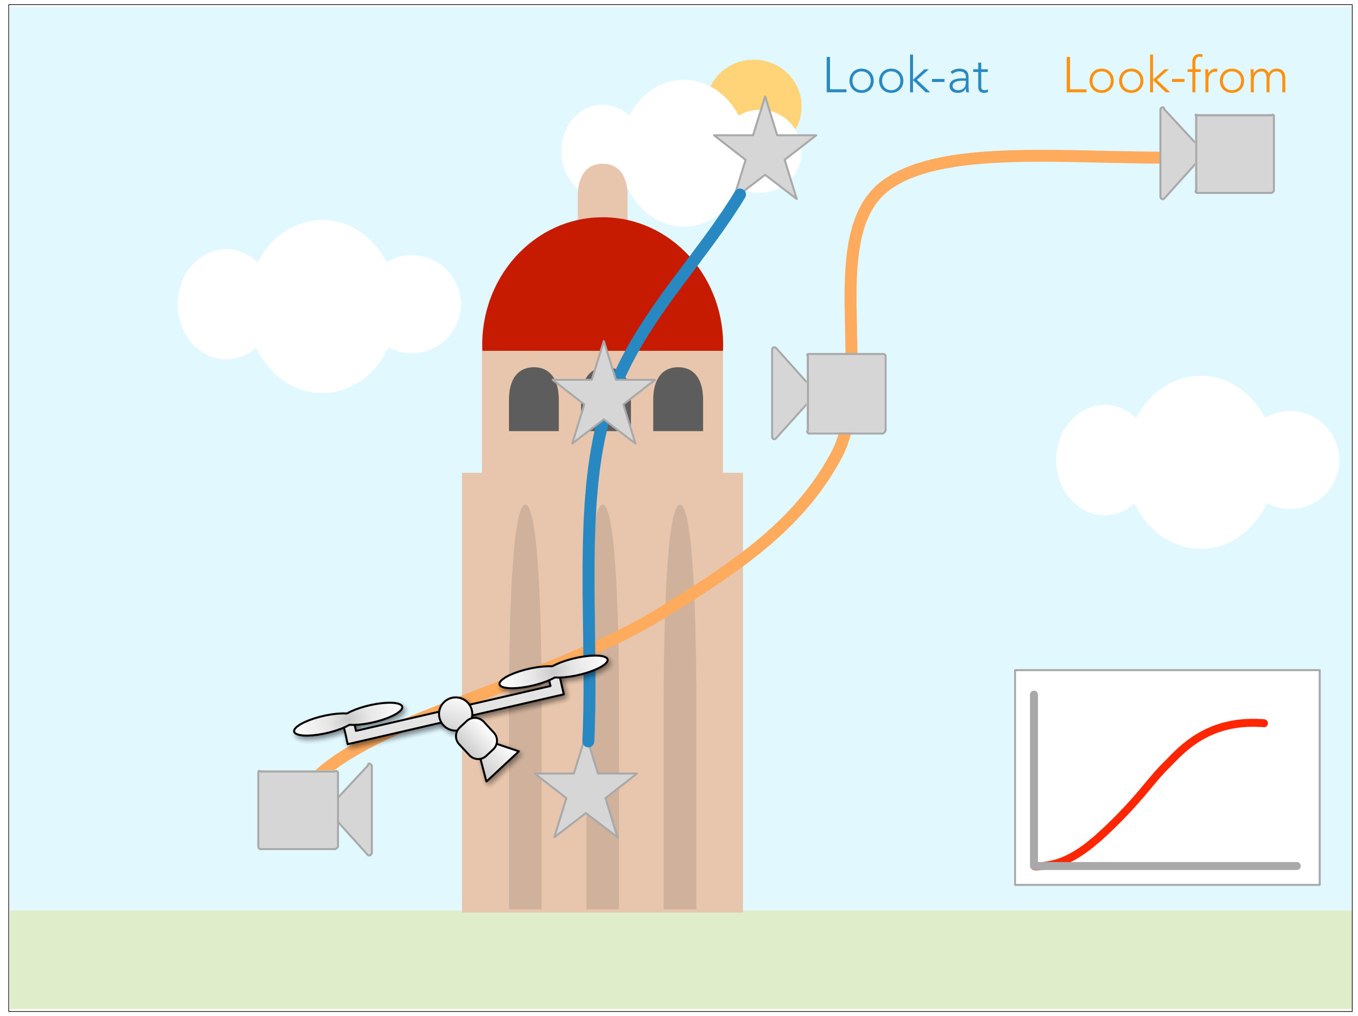

rotate and accelerate in a physically plausible way.

We describe how a quadrotor can rotate and accelerate using a set of equations called the equations of motion… So we can say the path has to

Satisfy Equations of Motion

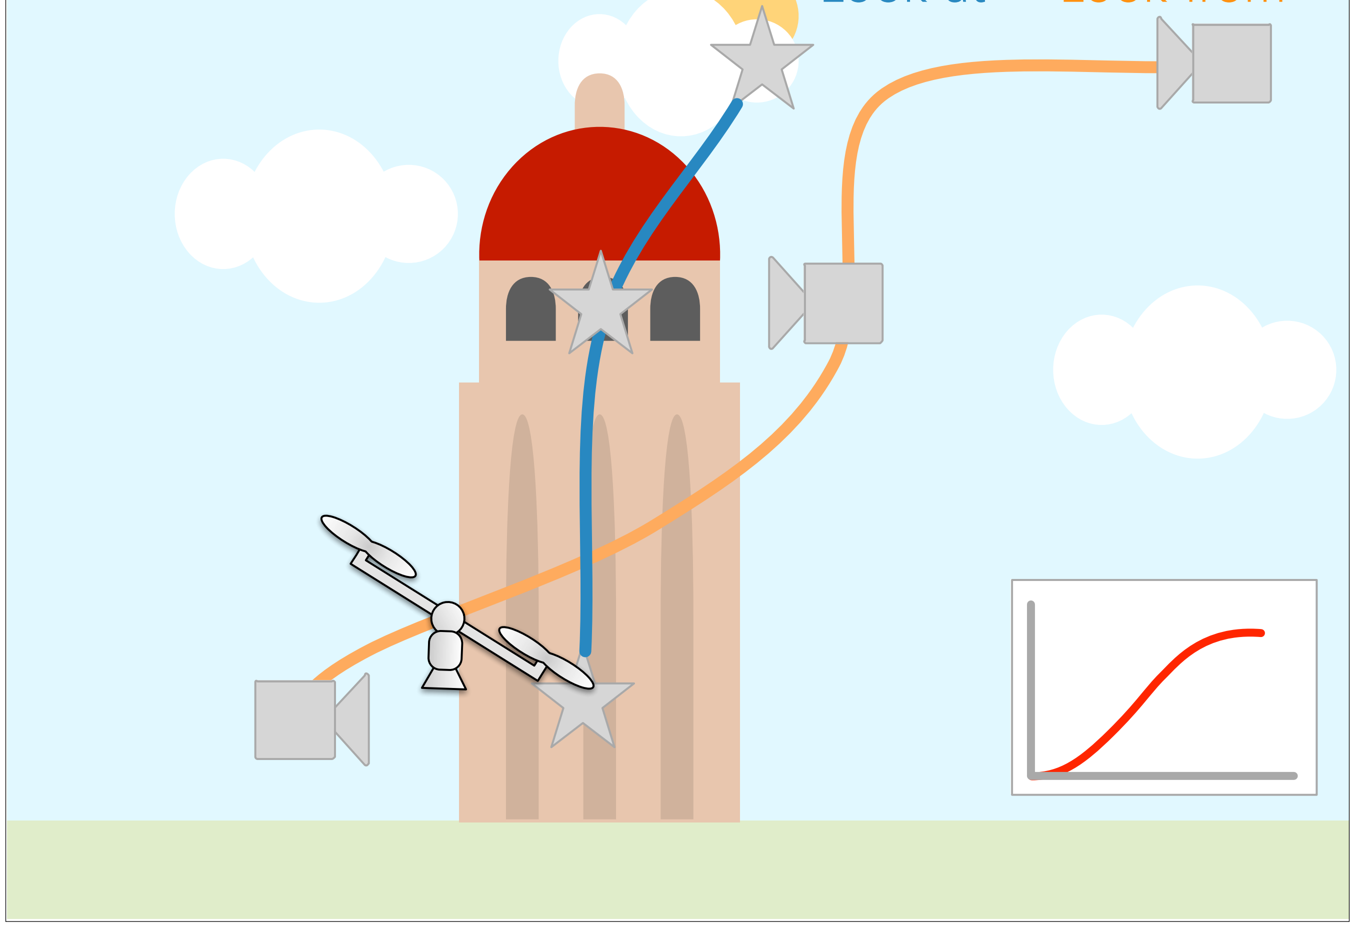

satisfy the equations of motion of our quadrotor camera

We also have to make sure the progress curve…

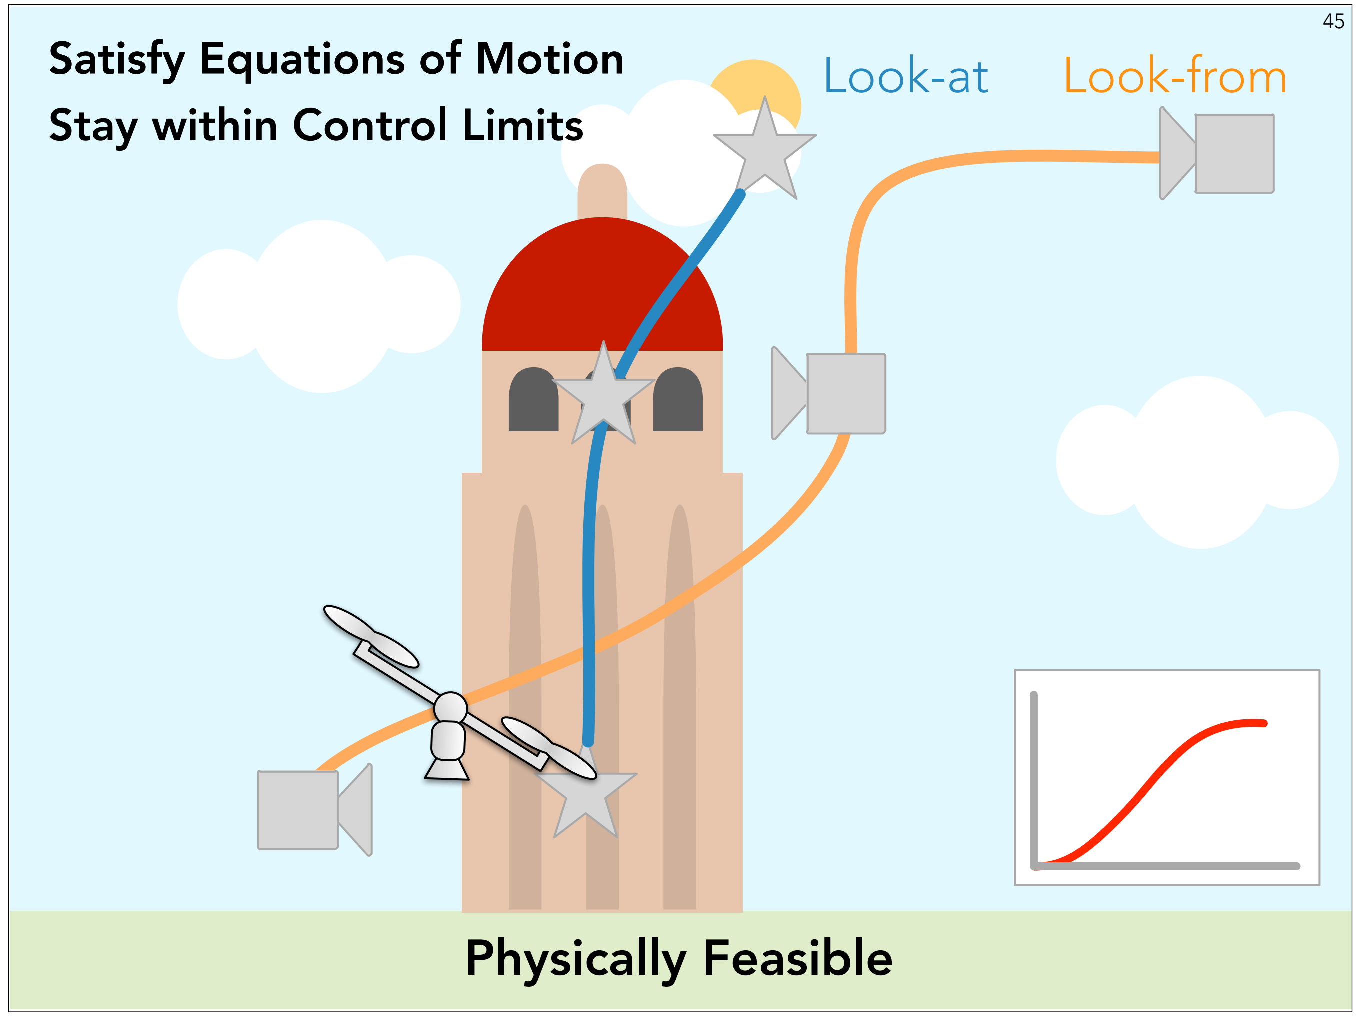

doesn't move the quadrotor too aggressively along this path.

We describe how aggressive a quadrotor can move by specifying control limits.

So the trajectory has to…

[CLICK] stay within our quadrotor's control limits.

If a trajectory satisfies both the equations of motion and the control limits,
[CLICK] we say it's physically feasible.

How do we generate physically feasible trajectories for quadrotor cameras?

CHALLENGE: HOW CAN WE FIND PHYSICALLY FEASIBLE TRAJECTORIES?

It's known that as long as the 4th derivate of a trajectory exits, then it **satisfies quadrotor equations of motion**.

[Mellinger et al, 2013]

How do we add a **camera** to this?

If we give users **timing control**,
how do we handle control limit violations?

It turns out we know quite a bit about the types of trajectories that quadrotors can follow. In fact, quadrotors have this surprising property.

[CLICK] As long as I can take the 4th derivative of whatever path I want to fly along, then it satisfies the equations of motion of the quadrotor.

[CLICK] But, we have more than just a quadrotor, we have a quadrotor camera. How do we add the camera, and how do we make that follow a look-at look-from path?

[CLICK] Secondly, if we give users fine-grained timing control, then they decide how aggressively we're flying. For example, they can set up shots that move really fast, faster than our quadrotor camera can. How do we handle control limit violations?

Let's start with adding the camera.

ADDING A CAMERA TO A QUADROTOR

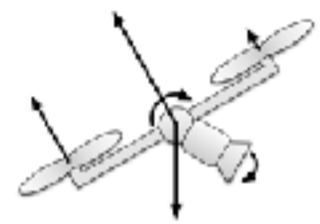

Physical model

We built a model of our quadrotor camera. This model describes the physical characteristics of quadrotor cameras.

Using this model, we proved that the same continuity requirement for a quadrotor's position path hold for both the look-from AND the look-at path! …

ADDING A CAMERA TO A QUADROTOR

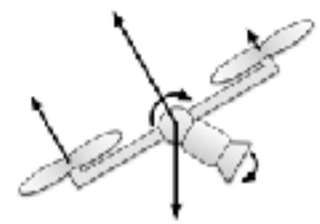

Physical model

Look-from and Look-at trajectory must both be C4 continuous.

Turns out that as long as the 4th derivative exists of the look-at trajectory and the look-form trajectory, then they satisfy the equations of motion of our quadrotor camera, and we can solve for how to fly and how to orient the camera. We call such curves “C4 continuous”.

For the sake of brevity, I won’t discuss the full proof, but if anyone is interested, I’ll refer them to our SIGGRAPH Asia 2015 paper.

GENERATING TRAJECTORIES FOR OUR QUADROTOR CAMERA

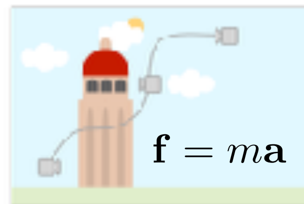

Generate trajectories that satisfy equations of motion

This is more continuity than is typically required in computer graphics applications.
For this reason, we couldn't use any off-the-shelf techniques to interpolate camera trajectories from the graphics literature.

To generate C4 continuous trajectories from our input keyframes…

GENERATING TRAJECTORIES FOR OUR QUADROTOR CAMERA

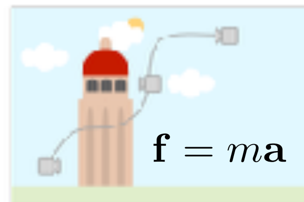

Generate trajectories that satisfy equations of motion

Interpolate using C4 continuous polynomial splines

Optimize for smoothness using a convex quadratic program

We use polynomial splines, we constrain them to provide C4 continuity, and we optimize the smoothness of our splines with a convex quadratic program over the spline coefficients. Again, I'll refer you to our 2015 Siggraph Asia paper if you're curious about the specific formulation.

The trajectories we generate here will satisfy the equations of motion given by our model,

GENERATING TRAJECTORIES FOR OUR QUADROTOR CAMERA

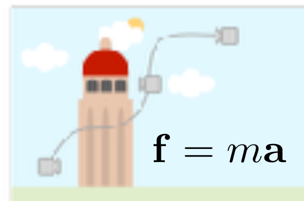

Generate trajectories that satisfy equations of motion **but not necessarily the control limits**

Interpolate using C4 continuous polynomial splines

Optimize for smoothness using a convex quadratic program

but won't necessarily be within the quadrotor's control limits.

The problem here is, the user has control over the timing of his shot, so he might give us a very aggressive progress curve that our quadrotor could never fly. As a silly example, If the user asks the quadrotor to fly from here to san francisco in 10 seconds, I can find a path, but I'm not going to be able to move along it fast enough.

DEALING WITH CONTROL LIMITS

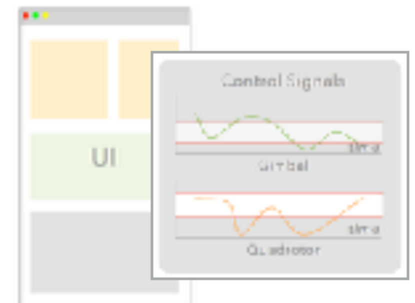

Display control forces in UI,
require user to stay within
control limits

We're going account for this by displaying the control forces for the trajectory in our user interface.
In case the user's trajectory exceeds any physical limits, the user can adjust their trajectory to be less aggressive.

TECHNICAL APPROACH FOR GENERATING FEASIBLE TRAJECTORIES

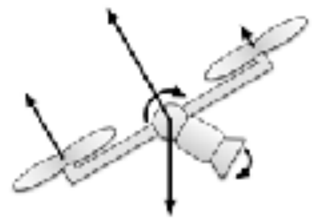

Physical model proves C4 continuity requirement holds for Quadrotor Cameras

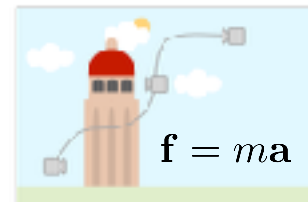

Generate trajectories that satisfy equations of motion but not necessarily control limits

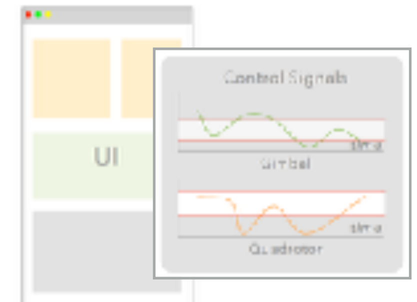

Display control forces in UI, require user to stay within control limits

To summarize,

We overcame the challenge of generating feasible trajectories by extending a previous proof for quadrotors to apply to quadrotor cameras.

We show how we can use the standard look-at look-from camera model to generate trajectories for quadrotors

And we show how we handle control limits, by providing visual feedback to the user.

IMPLICATION: EXPRESSIVENESS

We can express
any trajectory a quadrotor camera can execute

There's a really neat implication of this approach for our users:

At least theoretically, any path a quadrotor can fly, can now be expressed using keyframes and easing curves. We haven't given up expressivity moving from those manual control sticks to this formulation.

EVALUATION OF OUR TOOL

We built this tool and the underlying system, so let me show you that it actually works.

USER STUDY PROTOCOL

Four participants: **two novice** cinematographers, **two expert** cinematographers with quadrotor experience

90 minutes to create two feasible shots: one instructed, one open-ended

Autonomously capture shots on location

After capture, experts identified shot components that would make for challenging manual capture

We conducted an informal user study

with four participants. Two were novice cinematographers with no experience flying quadrotors. Two were professional cinematographers with extensive experience manually flying quadrotors to shoot video.

Each participant was tasked with creating two shots. For the first shot, we instructed them to capture a large belltower. For the second shot, we asked them to capture anything on Stanford's campus. We asked participants to create shots that they felt were challenging and impressive. We also asked them to make sure their shots are within the feasible limits of the quadrotor.

The participants then captured their shots fully autonomously.

After capture, we asked the experts to identify components in all of the shots that would be challenging to capture manually.

KEY QUESTIONS

How effective was our **visual preview?**

How did subjects use **feasibility feedback**?

Did we **enable cinematographers** to capture challenging shots?

From this study, I'm going to answer three questions:

first, "How effective was our visual preview".

second, "was feasibility feedback effective?".

Finally "did we enable cinematographers to capture challenging cinematography?"

VISUAL PREVIEW MATCHES FRAMING AND COMPOSITION OF CAPTURED SHOT

Expert 2, Freeform Shot

…our visual preview did indeed match the framing and composition of the captured shot well.

Here I'm showing the opening of our second expert's free-form shot. On top is the visual preview we generate, at the bottom is the actual video we autonomously captured.

Our system does have some failure cases.

<CLICK TO PLAY>

Here’s an example where the captured shot cuts off the top of a clocktower that’s visible in the rendered preview. This discrepancy comes from a few sources.

We depend on the accuracy of the Google Earth 3D model to plan our paths. Furthermore we depend on the accuracy of GPS-based quadrotor trajectory following. Sensor bias, noise and modeling inaccuracies can all cause discrepancies between the visual preview and the captured shot.

One possible approach we’re excited about is to use the rendered preview to visually correct the quadrotor’s flight in real time.

USERS CREATED FEASIBLE PATHS IN UNEXPECTED WAYS

We present the quadrotor's behavior on a set of feasibility plots, so the user can guarantee that the shot they create satisfies both their cinematic intent, and respects the hardware limitations of the quadrotor camera.

Some participants identified the feasibility plots as their favorite feature in our tool.

We examined the editing sessions to understand how our users ensured feasibility.

VISUAL FEASIBILITY FEEDBACK

All users successfully ensured feasibility

Complex behavior from feasibility feedback:

- Stretch the time
- Tweak trajectory
- Create completely different shot

<CLICK> All our users DID successfully create feasible shots.
<CLICK> Interestingly enough, users did this in several ways.
<CLICK> In some cases a participant simply slowed down their shot.
<CLICK> In others a participant chose to make major modifications to the visual content of the shot itself.
<CLICK> Tweaking the trajectory, or even
<CLICK> Creating a completely different shot.

I thought about this for a while, and I think there's an intuitive explanation. The user might be asking himself "hmm lets see if I can do this". If that's not possible, they might say "hmm let's try it slower and perhaps shorter" and come to the conclusion "this doesn't really get the mood I wanted across. Let's try something else"

Automating this process would be tricky, but I want to mention a cool extension to at least help with tweaking the timing of trajectories.

<CLICK>

ROBERTS AND HANRAHAN EXTENDED HORUS WITH AUTOMATED SMART TIME STRETCHING

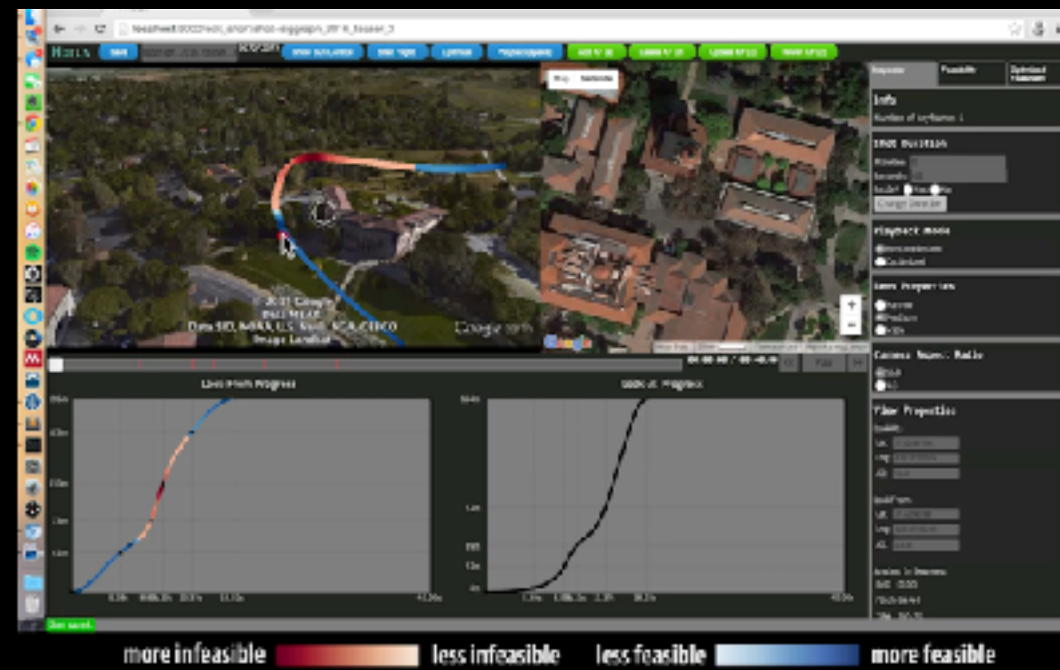

[Roberts and Hanrahan SIGGRAPH 2016]

Our collaborators Mike and Pat has since shown how to automate time-stretching shots in a way that gets as close as possible to the control limits of the quadrotor.

Their work is a great addition to further support rapid iteration on Horus. Since this is built right into our interactive tool, so if the user creates a shot that moves too fast, he can quickly see how close the quadrotor can get to his shot. A user can then choose to use the slowed-down shot, or decide to change his trajectory.

All built into the same system, which, btw, is open source and available online.

WE ENABLED NOVICE AND EXPERT CINEMATOGRAPHERS TO CAPTURE CHALLENGING SHOTS

I’d like to convince you that we enabled cinematographers to capture challenging shots

I’ll show you a variety of results our users produced. For each shot, we asked our experts to identify what would make it difficult to capture manually, which I’ll share.

REGARDLESS OF SKILL LEVEL, USERS CREATE CHALLENGING SHOTS

Expert 1, Instructed Shot

Here's the example shot created by our friend Weta Digital cinematographer.

Our expert explained that this shot is challenging since it requires continuously reorienting the quadrotor while simultaneously moving it along an arc. It's precisely timed so it looks directly at the tower just as it comes by, and it gets close enough that if you were flying this manually, you could easily crash into the tower. Capturing this with joysticks require a complex set of synchronized hand-movements.

REGARDLESS OF SKILL LEVEL, USERS CREATE CHALLENGING SHOTS

Novice 2, Instructed Shot

Here is the instructed shot created by someone who doesn’t know how to fly a quadrotor. Notice how its similar to the previous shot: the quadrotor also moves in an arc while continuously reorienting the camera to stay pointed at the tower.

REGARDLESS OF SKILL LEVEL, USERS CREATE CHALLENGING SHOTS

1.5X
Expert 1, Freeform Shot

1x
Novice 2, Freeform Shot

Here I'm showing you two more shots. On top is a shot from an expert. This shot is a very long single take revealing an environment. Such long takes are notoriously difficult to capture manually, since any small momentary mistake ruins the entire shot. On the bottom is a shot from a novice. This shot moves through a scene full of obstacles, where a small mistake by the pilot can easily cause a crash.

Using our system, both novice and expert users captured a variety of challenging cinematography.

I then went and designed a few shots myself to push the tool.

Here’s a first one that combines multiple of the aspects our experts thought were challenging. It’s long, it’s close to obstacles, it has some speed variation in it, and the camera rotates and moves in a complex way.

Here's a nighttime timelapsed shot we set up with Horus. Long and slow shot. If i was flying this by hand, I couldn't make any course corrections half-way, that would show up as a jerk. This runs at 10x normal speed, it amplifies any motion, so even if I think I'm being very smooth on the controls, I might still not get the shot, and I won't know until I edit my video to produce this.

I like to stare at this one. It reminds me of the movie Tron.

SUMMARY

Demonstrated a tool in which cinematographers could express high quality shots for quadrotor cameras, even if they were novices.

Showed how 3D Animation primitives and the look-at look-from camera model can be used to fly quadrotor cameras.

So with Horus

<CLICK> We….

<CLICK> Showed ….

WITH HORUS, THE USER ACTS AS A **CINEMATOGRAPHER**

COMPOSING SHOTS USING KEYFRAMES AND EASING CURVES

With Horus, the user acts as a cinematographer. They are only concerned with visual composition of shots over time, which they express using keyframes and easing curves. We raised the level of abstraction for quadrotor control by providing cinematography-first controls, at least for all cases where a user is okay with preplanning a shot.

PROBLEM

QUADROTOR CINEMATOGRAPHY REQUIRES

TECHNICAL SKILL OF FLYING, AND ARTISTIC SKILL OF COMPOSITION, SIMULTANEOUSLY APPLIED IN REAL TIME

In terms of our original problem statement, we removed two big requirements.

PROBLEM

QUADROTOR CINEMATOGRAPHY REQUIRES

~~TECHNICAL SKILL OF FLYING~~, AND
ARTISTIC SKILL OF COMPOSITION,
~~SIMULTANEOUSLY APPLIED IN REAL TIME~~

by **preplanning** camera movements

We no longer need the technical skill of flying, or to make real-time decisions.

Of course, they still need to have some artistic skill. This tool leaves it up to the user to set up interesting compositions. And yes, you can create the crappy shots with it. They'll be nice and smooth and feasible, but they can be badly composed.

And of course, there are some scenes for which preplanning a shot isn't appropriate, like, scenes in which you're filming a person, and you're not exactly sure where they'll be or what they'll do.

OVERVIEW

Compose shots using classic 3D Animation primitives, adapted to respect quadrotor camera physics [SIGASIA 2015]

Horus
A Tool for Shot Planning

Compose shots in real time using visual composition principles from filmmaking [arXiv 2016]

Drone Cinematographer
A Tool for Filming People

Now that we can capture cinematography by preplanning our compositions,

I'll move onto the next step.

We're going to take a stab at making the intricacies of visual composition more accessible to novices. And we'll look at this in the context of something our tool isn't great at - filming people, where you might want to make decisions in real time about camera placement, without the burden of manually flying a quadrotor. So we'll build a real-time system that doesn't require as much preplanning.

This is work I did in collaboration with Jane, Dan, Floraine, Mike and Pat

SHOT WITH DRONE CINEMATOGRAPHER

And here’s a shot captured using the Drone Cinematographer. This time there’s no professional cinematographer either, just me using a simple tap-to-fly interface.

Pretty neat huh!

HOW MIGHT WE AUTOMATE THE COMPOSING OF SHOTS?

How might we automate the composition of shots?

To do this, we'll start by investigating visual composition…

WHAT DO WE KNOW ABOUT VISUAL COMPOSITION? ABOUT SHOTS?

Turns out, there’s a whole lot of expert knowledge in how to compose shots.

There exists these guiding principles for visual composition, and there’s a whole taxonomy of different types of shots. This is all captured in the language of film developed by traditional cinematographers over the last century.

Let me give you a few examples of the visual composition principles that we might use for shot planning…

EXAMPLE OF VISUAL COMPOSITION PRINCIPLE: RULE OF THIRDS

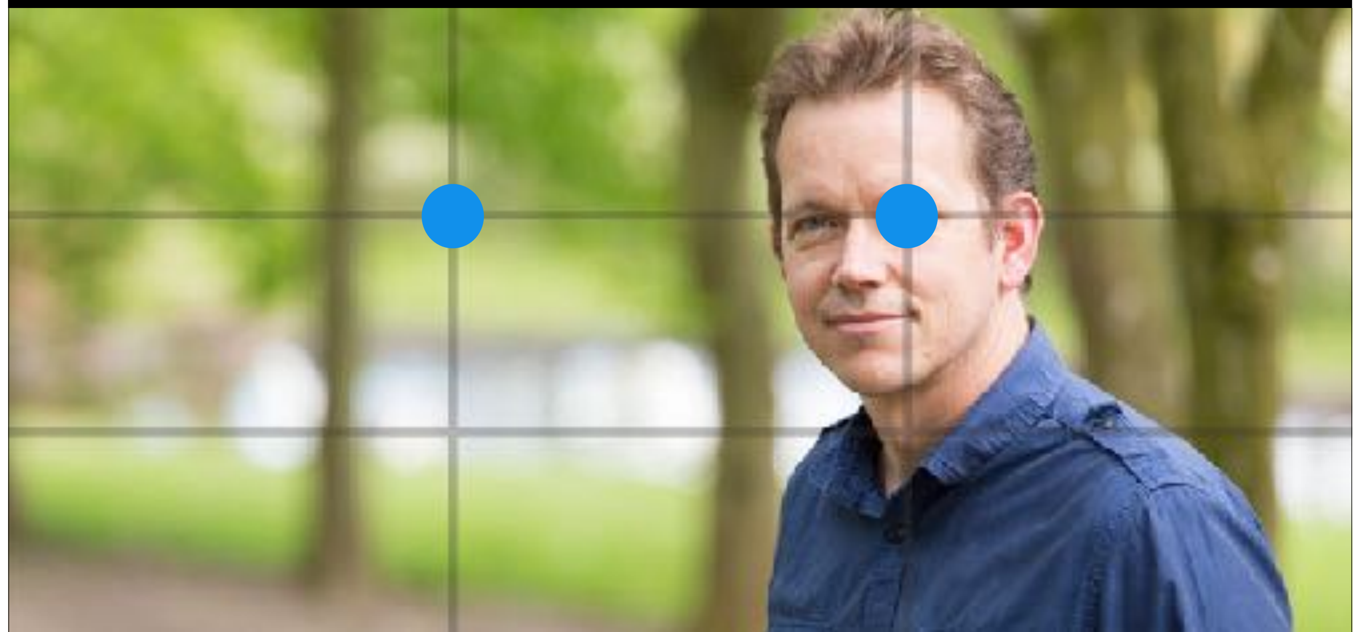

[Ryan Bromley, 2016]

Here's an example of a visual composition principle that many of you will recognize: the rule of thirds. This is really a collection of principles depending on your style and what you're capturing, so I'll give you two examples.

In case of capturing subjects, it states that subjects should be placed at these attractor points [CLICK] in blue.

EXAMPLE OF VISUAL COMPOSITION PRINCIPLE: RULE OF THIRDS

[DesignShovel, 2016]

the rule of thirds also suggests which side of the frame to place the subject. if they're facing to the right, place them on frame left. This creates a sense of space into which they can move.

FILM LANGUAGE DESCRIBES CANONICAL SHOTS

[Arijon, 1976]

Film language also describes a set of canonical shots that comes up all the time in traditional filmmaking. I'll show you two simple ones here, but there are more.

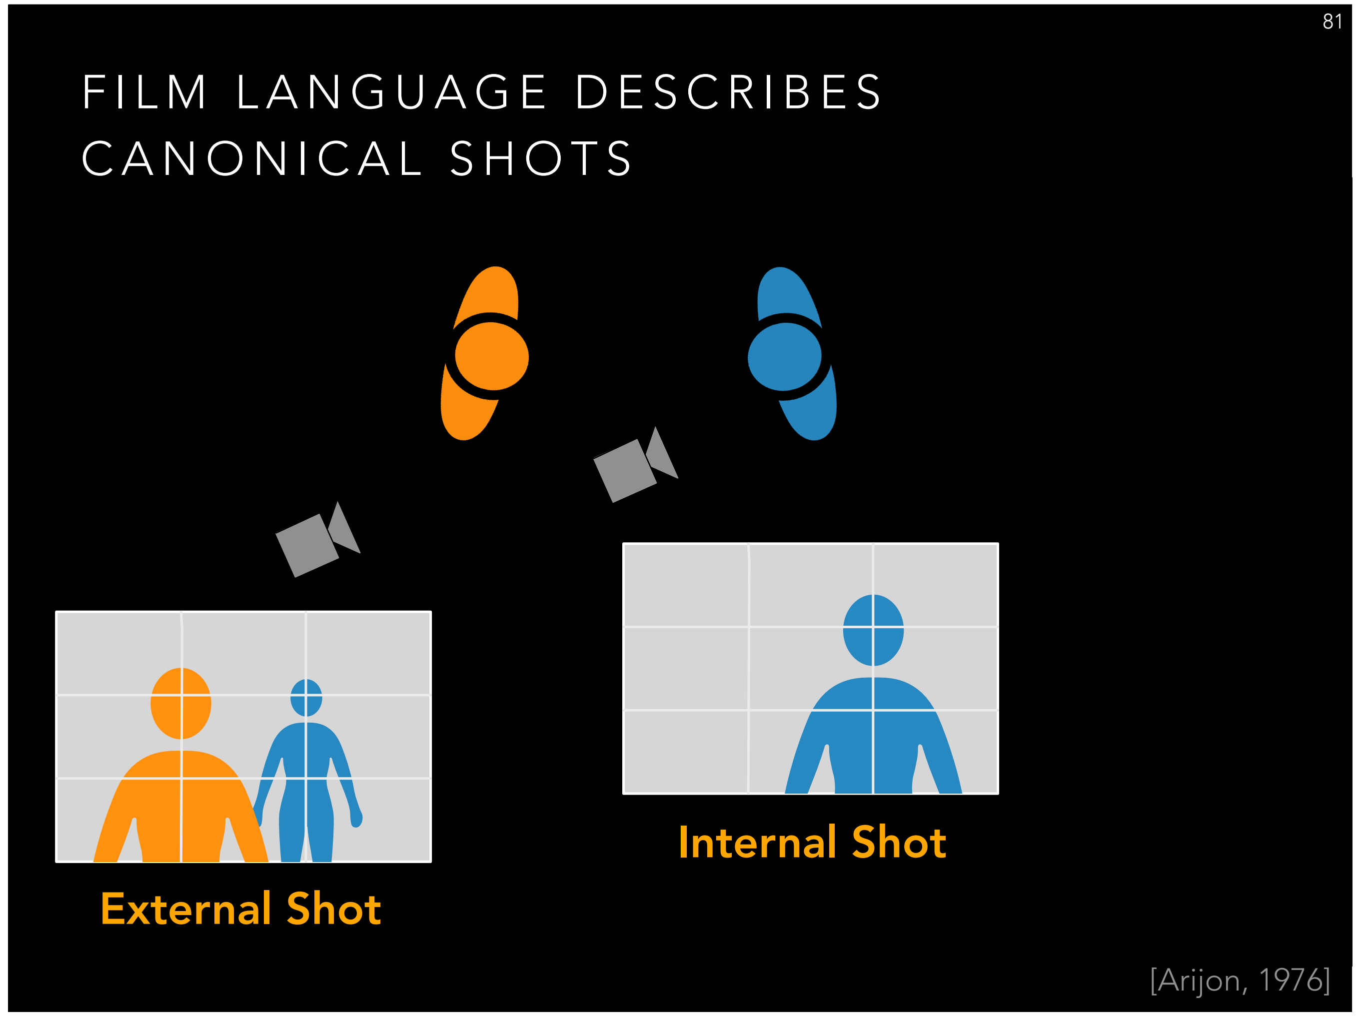

For example, the External shot, where one person is filmed over the shoulder of another person.

and the Internal shot, which is a closeup of one person.

I do want to point out that film language has a very rich body of *static* camera locations: places where the camera is stationary on the ground, or stationary relative to people. (eg. in back of car) In general it champions “lazy cameras” - in most cases the camera moves much slower than the action of the scene.

Let me show you another visual composition principle…

EXAMPLE OF VISUAL COMPOSITION PRINCIPLE: RESPECT LINE OF ACTION

Here’s a principle that informs how we can move cameras so subsequent shots sequence together nicely. This principle is called the “180 degree rule” or “respecting the line of action’

EXAMPLE OF VISUAL COMPOSITION PRINCIPLE: RESPECT LINE OF ACTION

[Arijon, 1976]

Let's say I'm capturing this external shot, looking a the blue person from behind the orange person.

Next i want to move the camera so we can see the face of the orange person. Where's a good place to set up another external shot…

EXAMPLE OF VISUAL COMPOSITION PRINCIPLE: RESPECT LINE OF ACTION

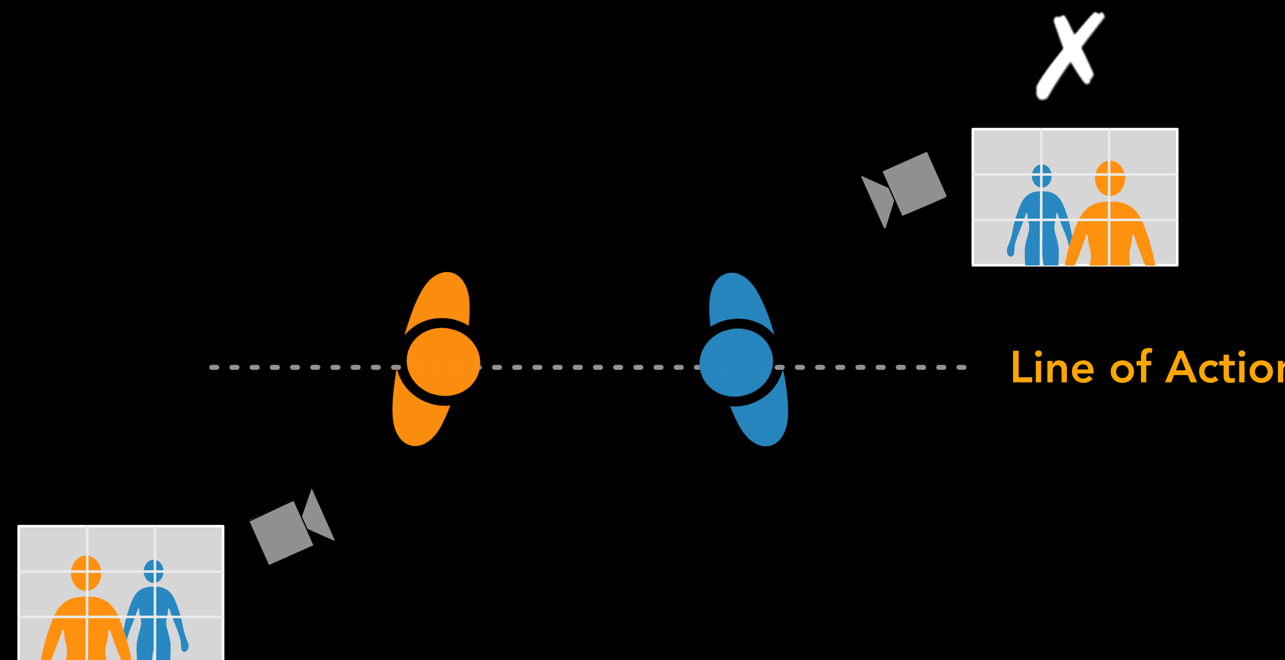

[Arijon, 1976]

If I cross the line of action, something funny happens - the two people swap places in the frame. This is jarring…

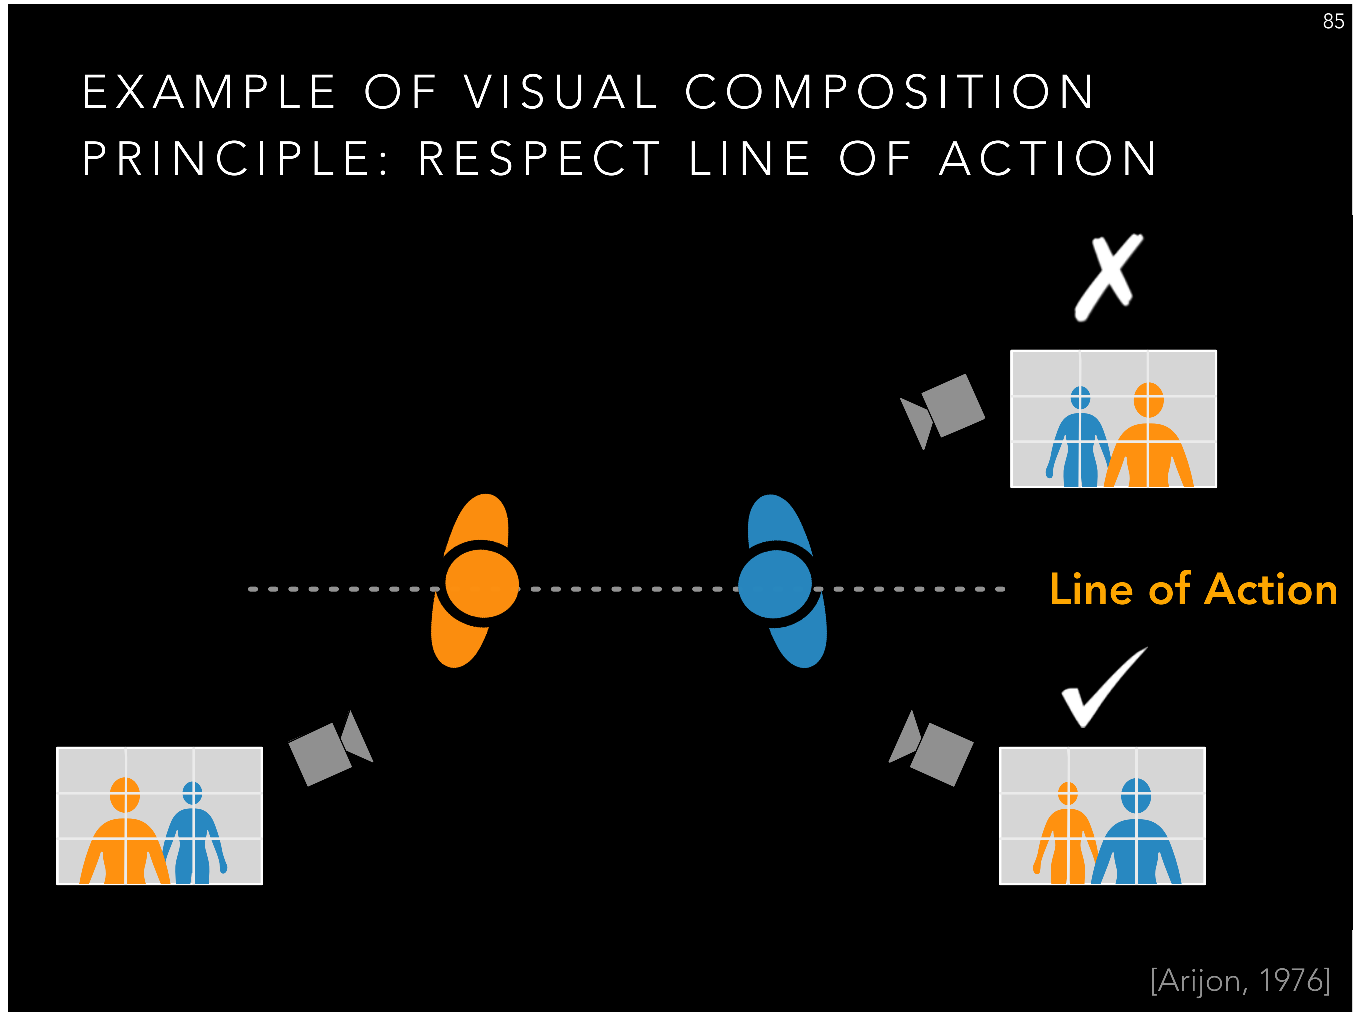

while if we stay on the same side of the line of action, the two people stay on the same sides of the frame. If we cut between these two shots, we won't get a surprising flip of characters on screen.

THESE IDEAS HAVE BEEN USED IN COMPUTER GRAPHICS TO CONTROL CAMERAS

[He et al, 1999]

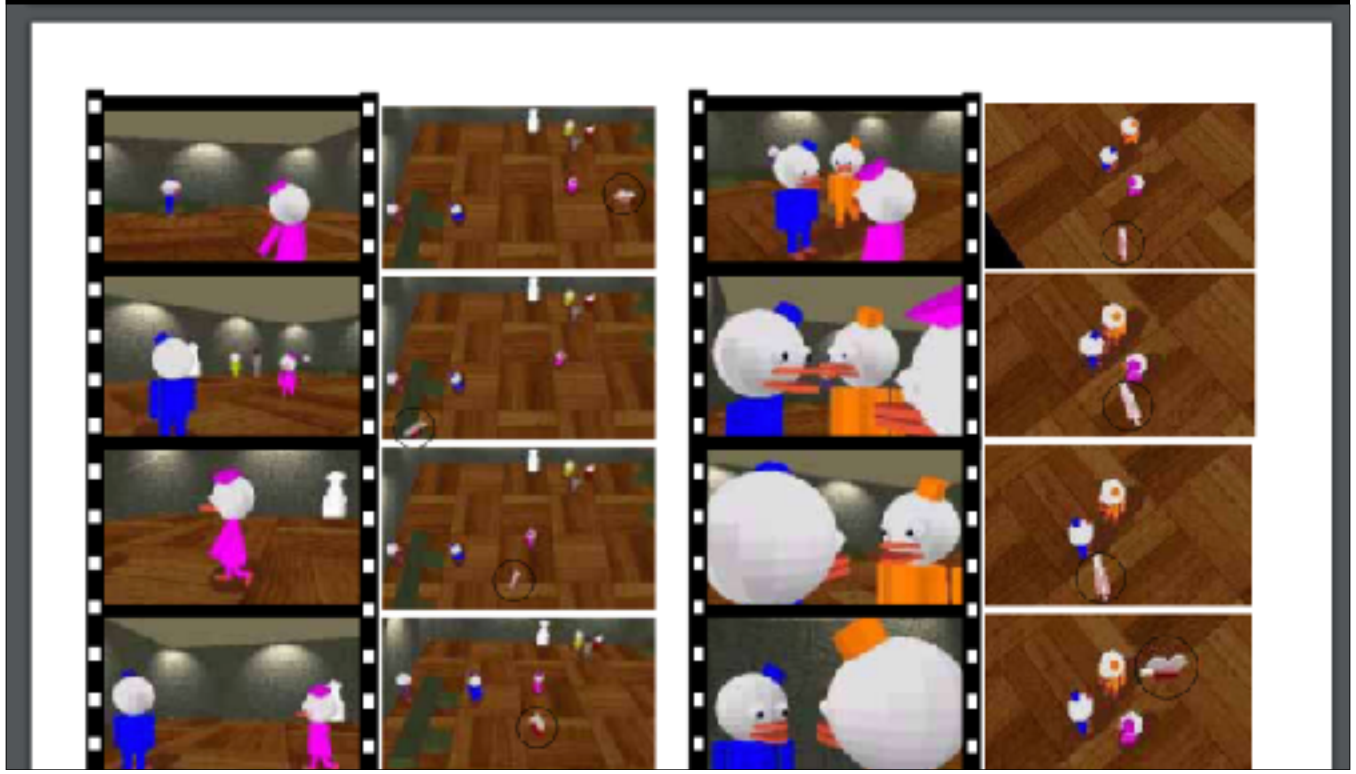

The computer graphics community have used these ideas to automate virtual cameras! Once again, we need to just do one weird trick… and adapt these methods to the demands of the physical world…

You'll notice a pattern here.

Armed with a set of visual composition principles and canonical shots from this paper and the previous books i mentioned…

KEY IDEA

Use **canonical shots and composition principles** as control input to a Quadrotor Camera

Automate flying the quadrotor

Novices get **high quality visual composition**

Can **react in real time**

We can build a tool that lets the user fly their quadrotor by picking one of these canonical shots.

[CLICK] We'll automate flying the quadrotor camera.

[CLICK] We'll make sure our shots, and the quadrotor's movement between shots, follow visual composition principles. Now novices can get high quality visual composition for free.

[CLICK] We'll do this in real Time, so we can react to what our subjects are up to.

Let me show you the canonical shots we expose to users…

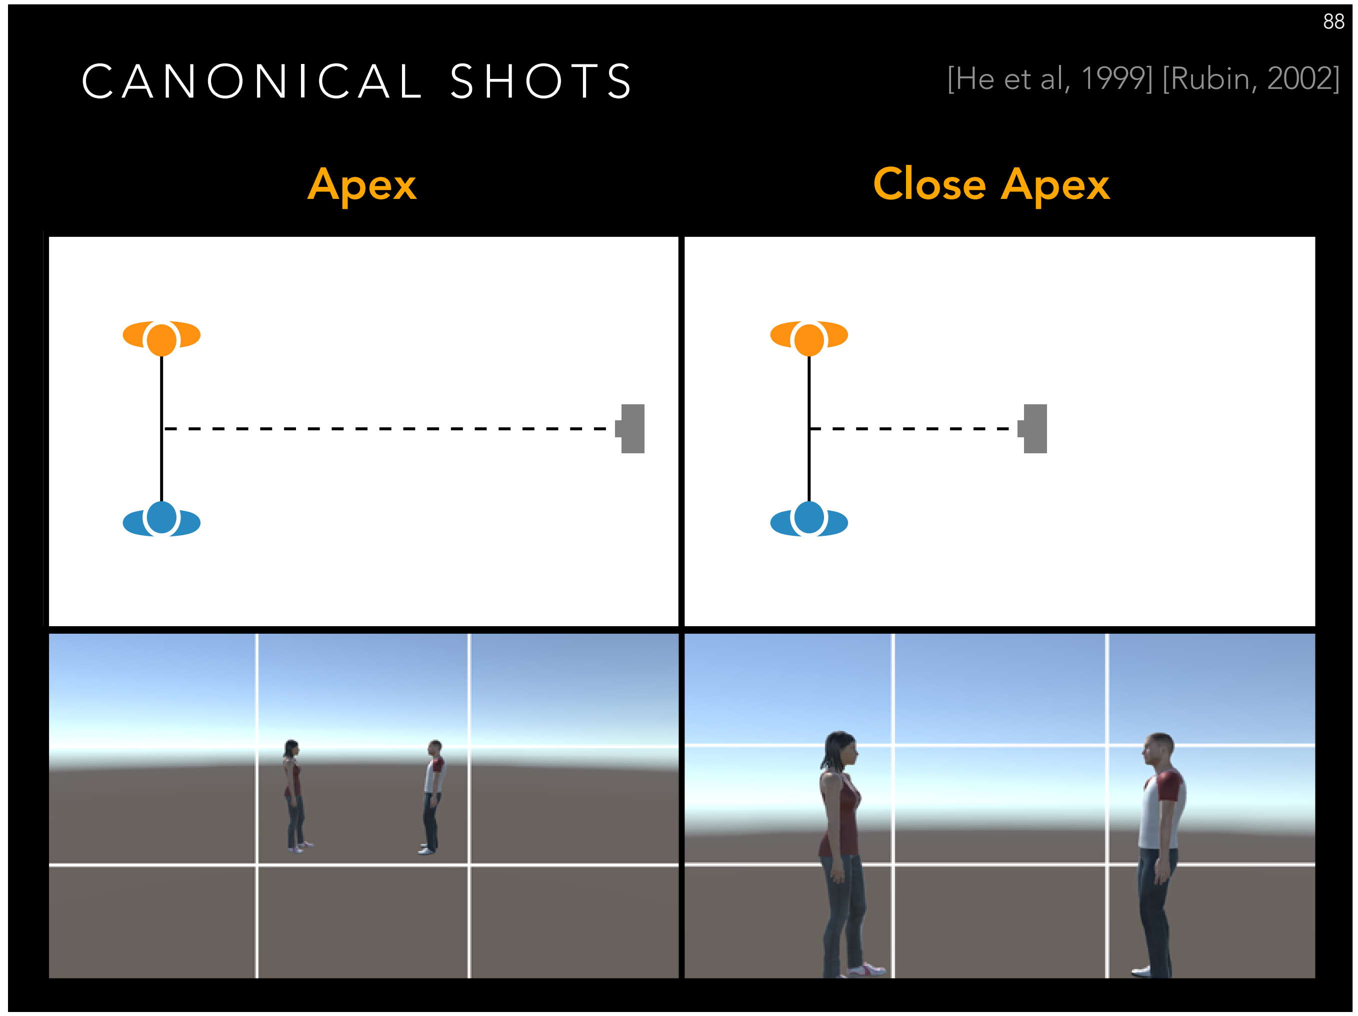

(Talk through shots)

Let me show you the canonical shots we expose in our system.

For each of these shots, we have an equivalent shot looking from a higher vantage point.

We adapted these shots from He et al's Virtual Cinematographer, who investigated using canonical shots to drive cameras in virtual reality environments, We also took inspiration from Michael Rubin's 2002 book on Digital Video.

With a fairly small set of shots we can capture a lot of interesting stuff.

We can add more shots easily, since they're defined by the position of the actors on screen, the angle of the camera, and the distance to the line of action.

CANONICAL SHOTS

[He et al, 1999] [Rubin, 2002]

For every subject…

(Talk through shots)

Let me show you the canonical shots we expose in our system.

For each of these shots, we have an equivalent shot looking from a higher vantage point.

We adapted these shots from He et al's Virtual Cinematographer, who investigated using canonical shots to drive cameras in virtual reality environments, We also took inspiration from Michael Rubin's 2002 book on Digital Video.

With a fairly small set of shots we can capture a lot of interesting stuff.

We can add more shots easily, since they're defined by the position of the actors on screen, the angle of the camera, and the distance to the line of action.

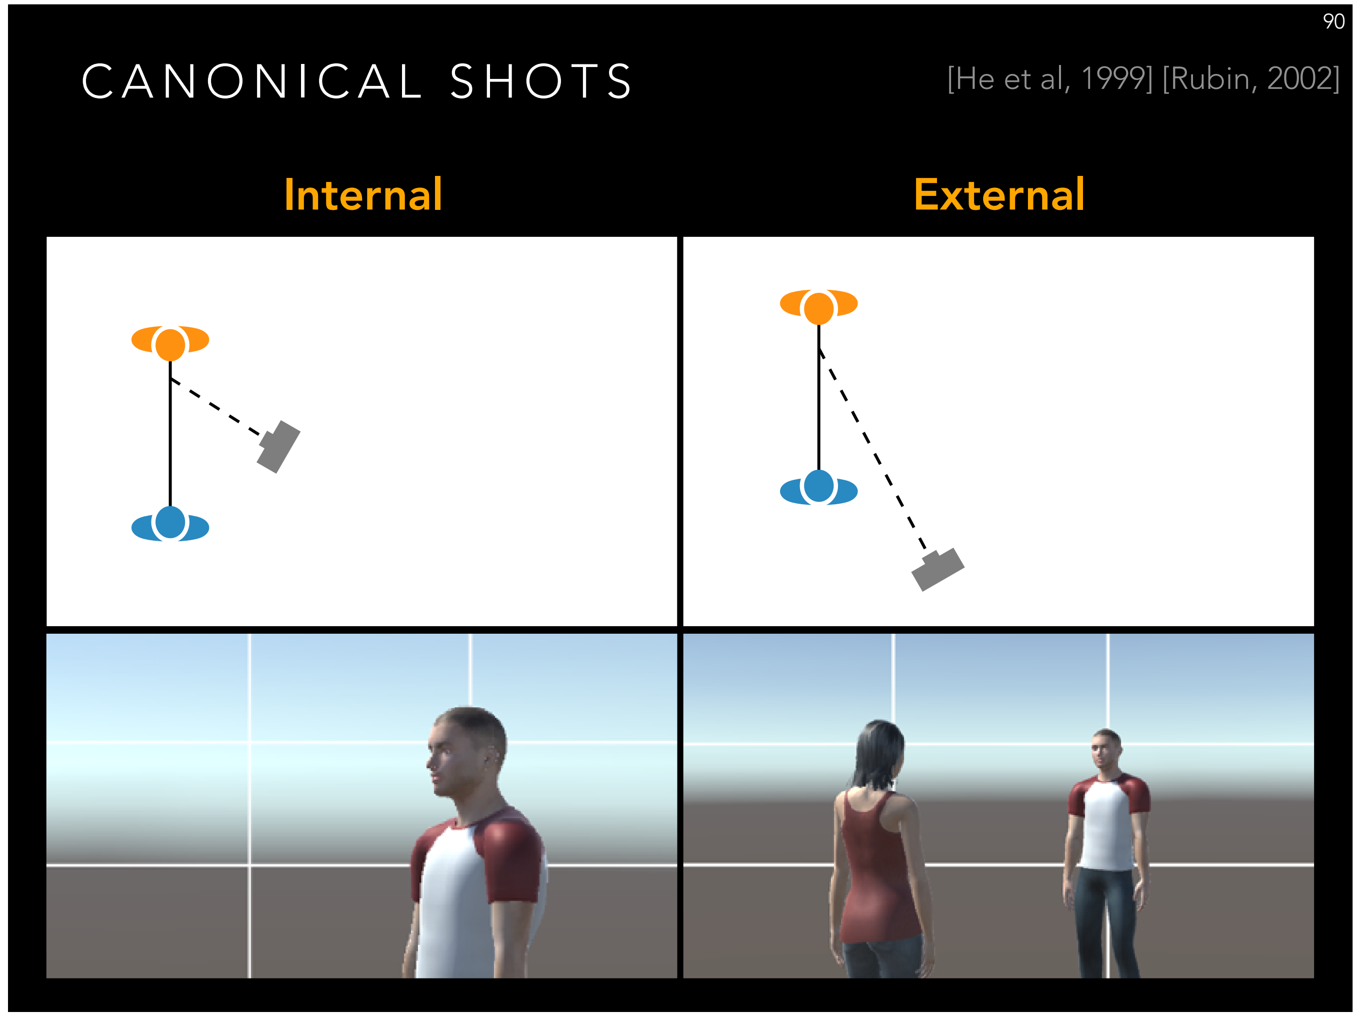

(Talk through shots)

Let me show you the canonical shots we expose in our system.

For each of these shots, we have an equivalent shot looking from a higher vantage point.

We adapted these shots from He et al's Virtual Cinematographer, who investigated using canonical shots to drive cameras in virtual reality environments, We also took inspiration from Michael Rubin's 2002 book on Digital Video.

With a fairly small set of shots we can capture a lot of interesting stuff.

We can add more shots easily, since they're defined by the position of the actors on screen, the angle of the camera, and the distance to the line of action.

CANONICAL SHOTS

[He et al, 1999] [Rubin, 2002]

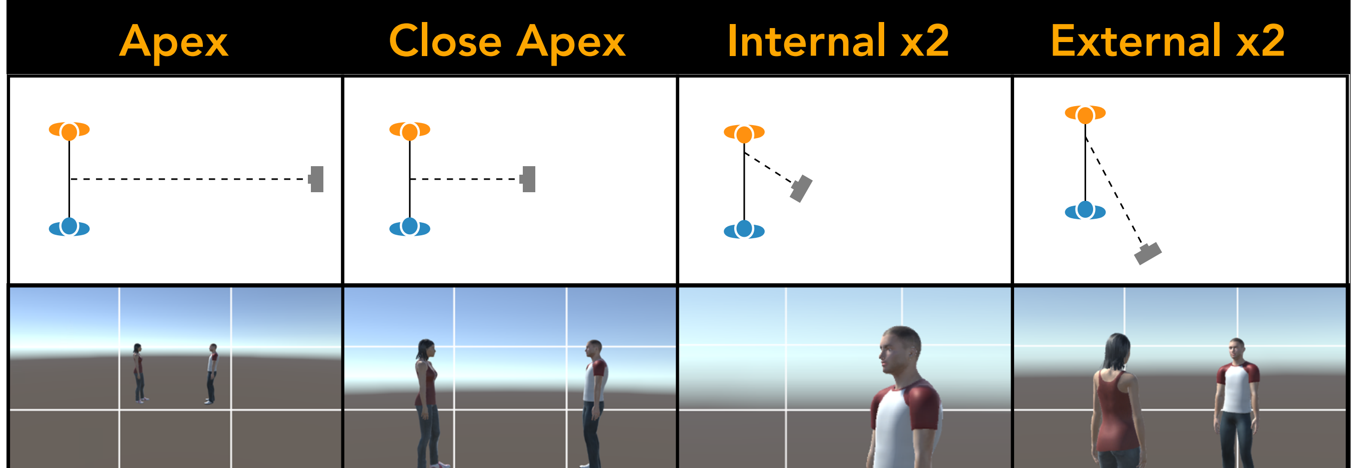

And the equivalent shots from a **higher vantage point**

For each of these, we also have the equivalent…

from a higher vantage point

USER INTERACTION

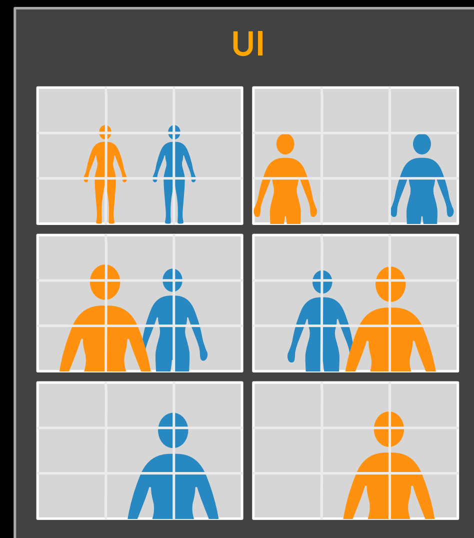

This list of canonical shots is exposed in a user interface.

USER INTERACTION

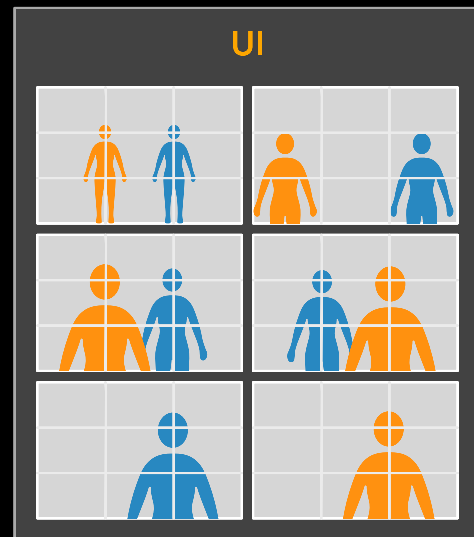

If we want to film our two subjects.

USER INTERACTION

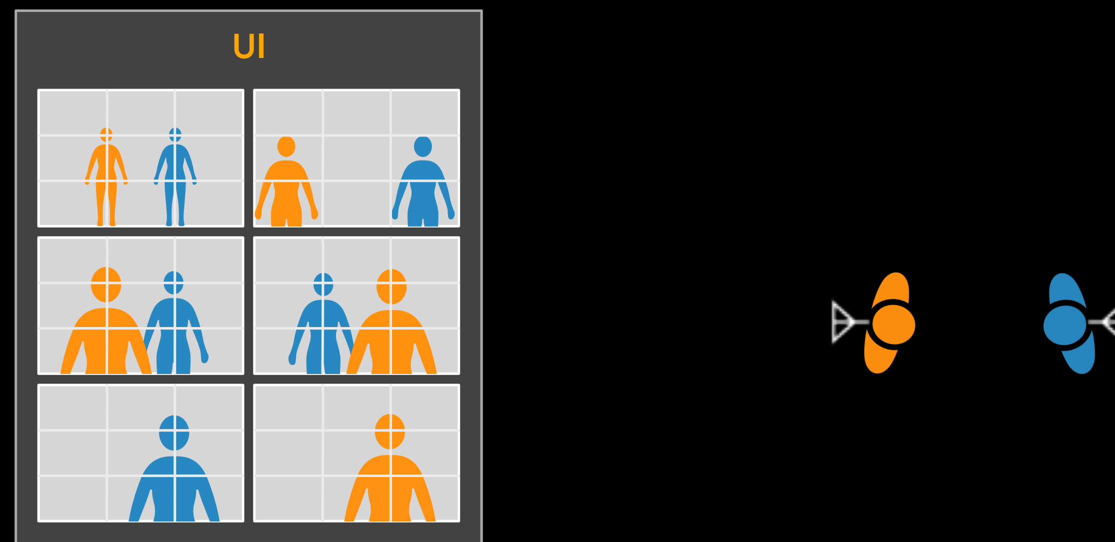

We need to know where they are, so we place trackers on each subject

USER INTERACTION

When the user clicks on a canonical shot in our interface…

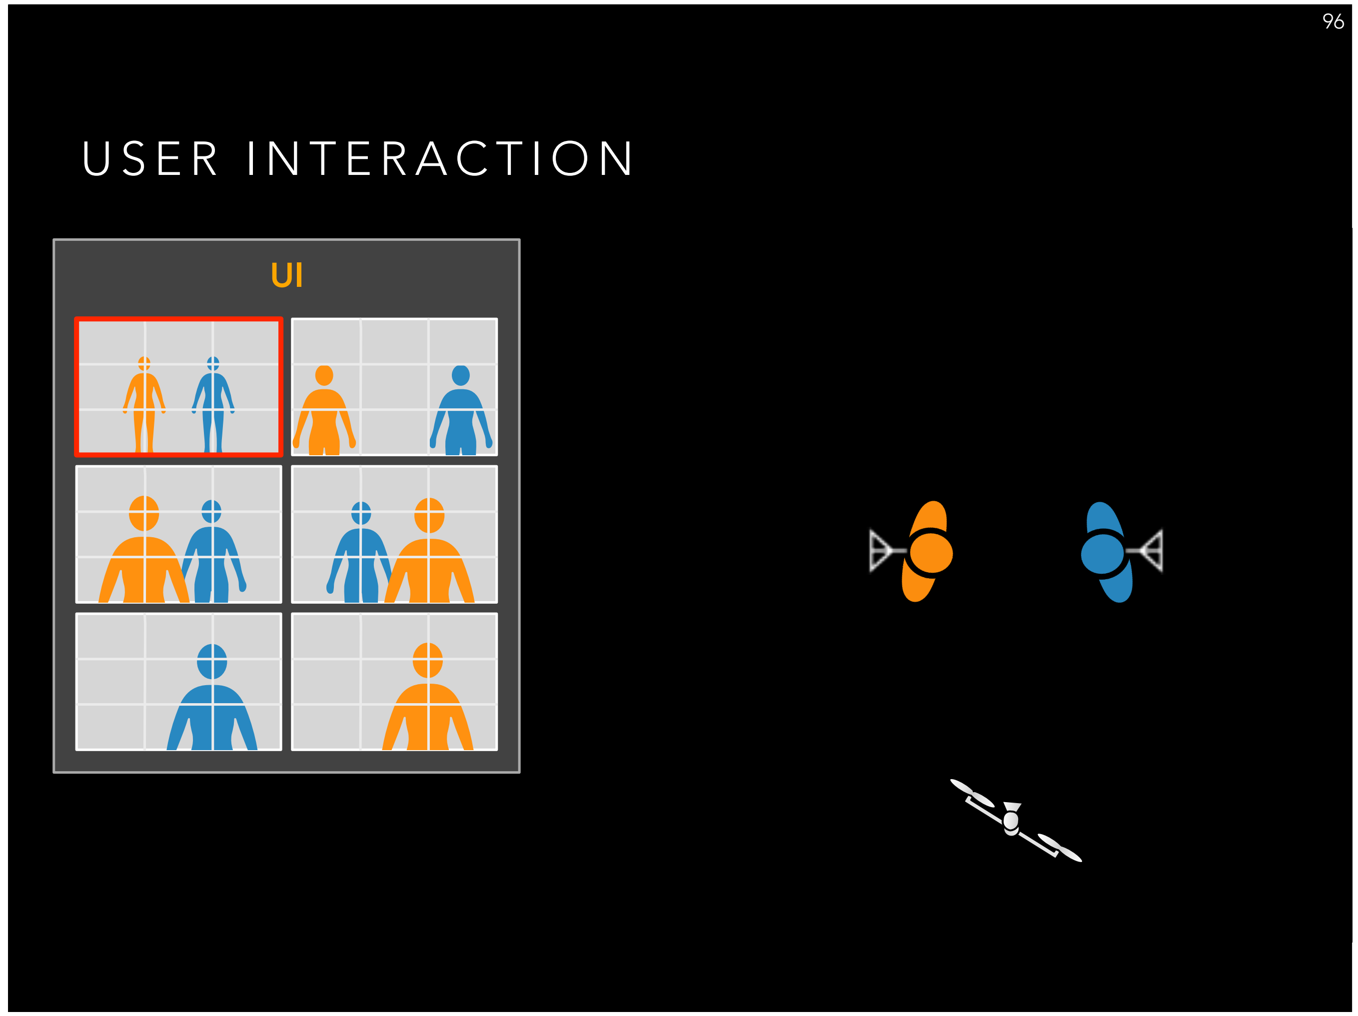

… our system will automatically capture it with a quadrotor. The quadrotor will keep hovering, holding this position, mimicking a lazy camera.

TOOL OVERVIEW

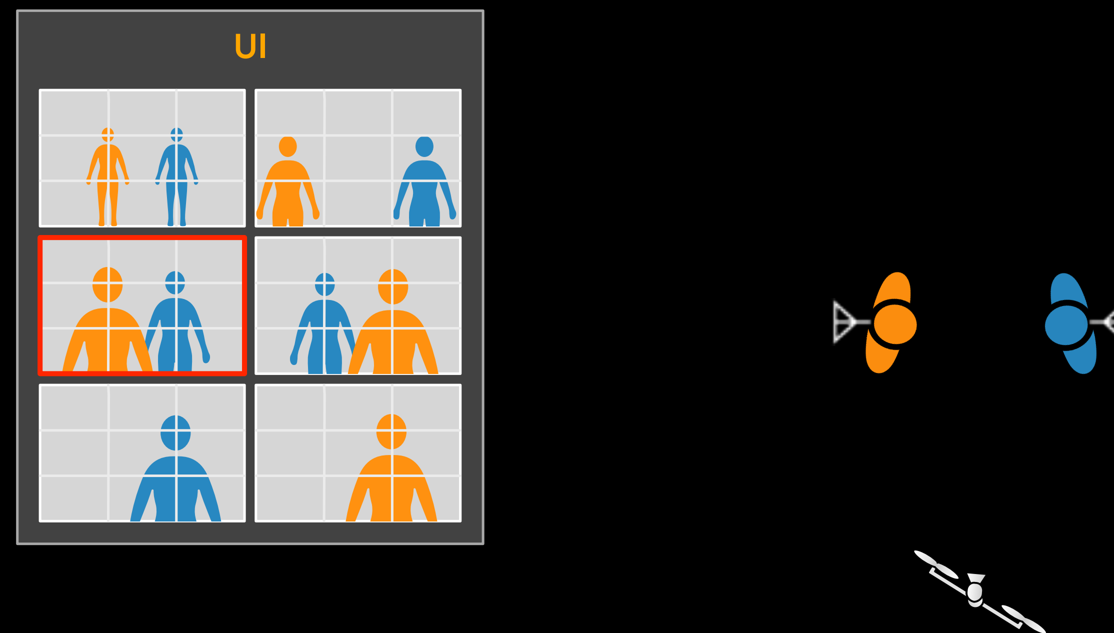

Whenever the user clicks a next shot

TOOL OVERVIEW

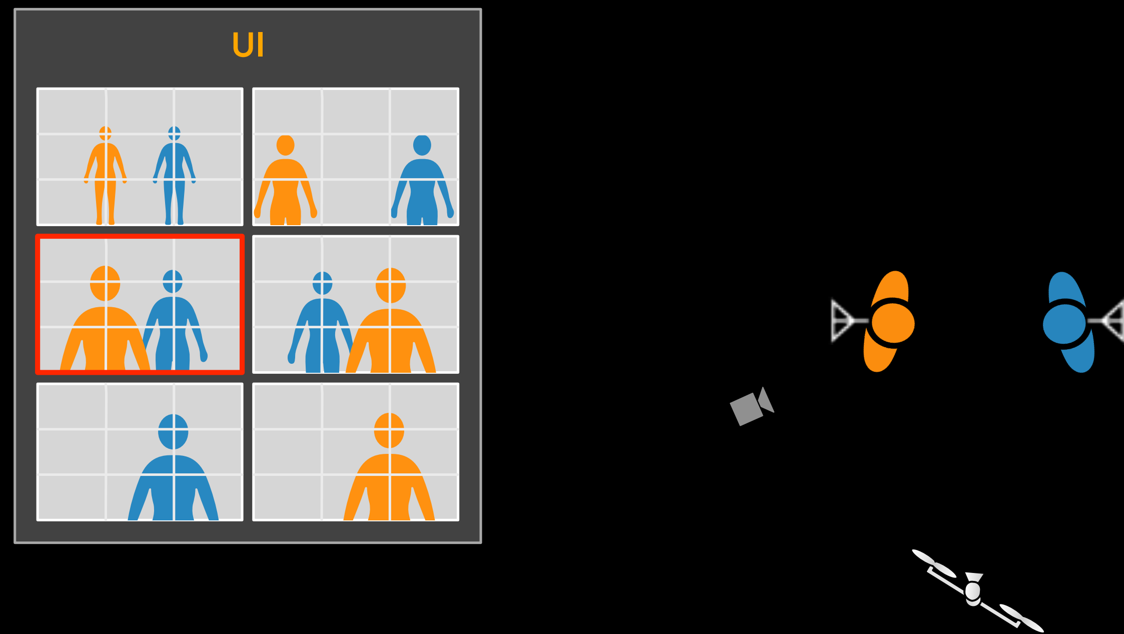

our system calculates the camera pose to capture this shot using the rules of visual composition and the current position of the subjects

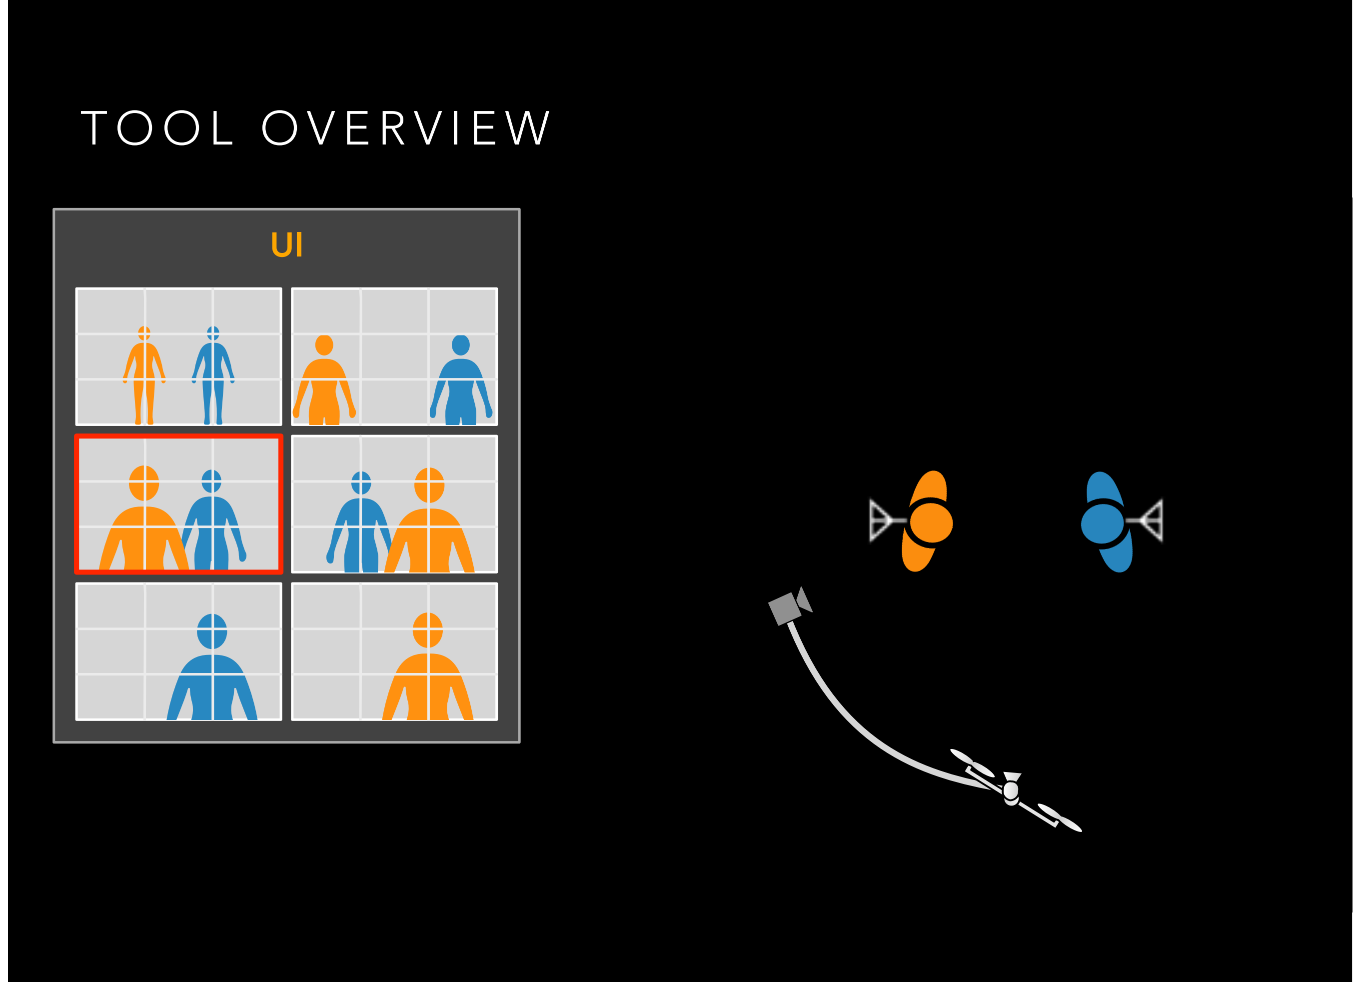

we find a high quality trajectory to get to this camera pose. Because this is all happening in real time, we want to capture high quality video along this trajectory, so we take extra care to make sure this trajectory also follows some of our visual composition principles, and what we know about creating feasible trajectories…

Once we have this trajectory…

TOOL OVERVIEW

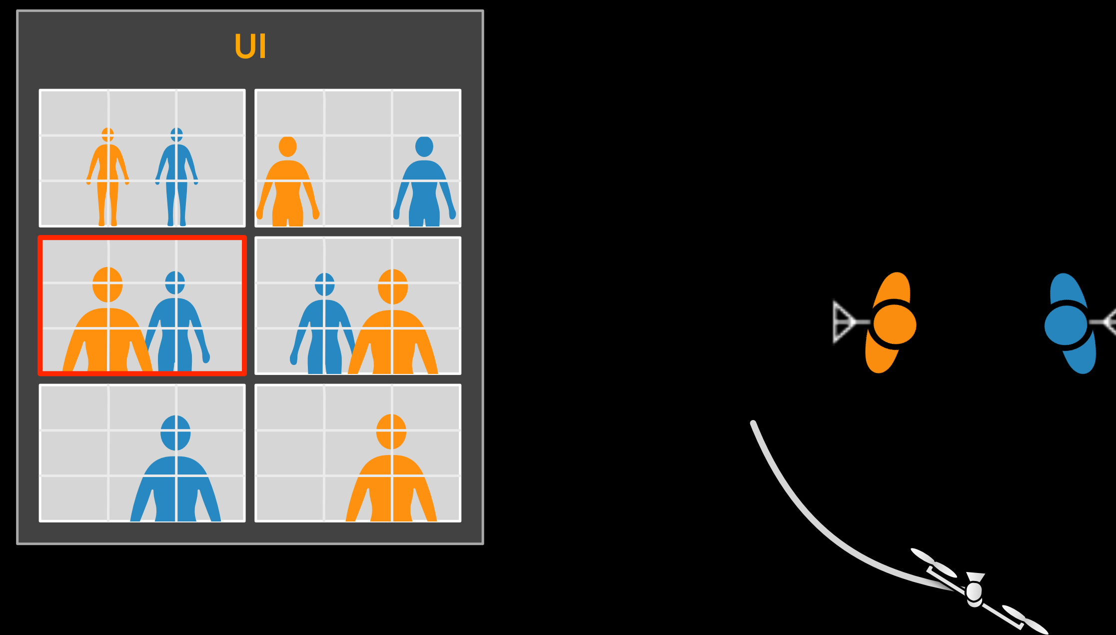

… we automatically fly the quadrotor to this new position.

To build this system we have to overcome a set of…

CHALLENGES

We need to **track** our subjects

Keep our subjects **safe**

Find **high quality transitions** between shots

Technical challenges. I'll tell you about three of them, and they're all documented in our paper…

First, since we want to film people, we need to [CLICK] track them, and it turns out off the shelf methods don't quite cut it.

Second, we want to [CLICK] keep our subjects safe, so we have to be careful where we place and how we move the quadrotor.

Lastly, we need to [CLICK] have a way to find high quality transitions between these shots.

SAFETY

Let me first talk about safety, since this is a particularly important point when filming people with quadrotor cameras.

SAFETY

RESPECT A NO-FLY SPHERE

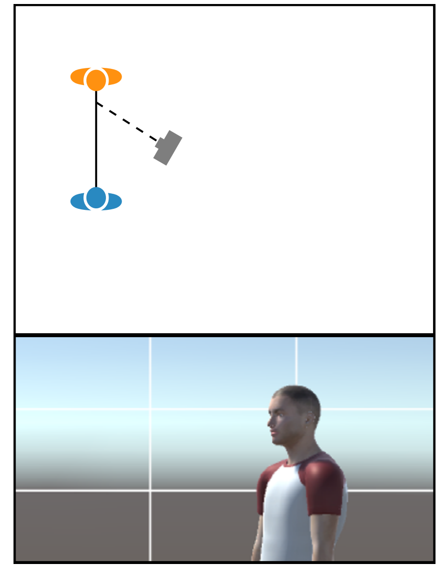

We'll approach this problem by defining a safety sphere around every person, indicated in green. We calibrated this safety sphere during flight to a value our users felt comfortable with.

Currently quadrotor cameras are mostly equipped with fairly wide angle fixed lens cameras. So for many of our shots, the default camera placement is quite close, and inside this safety sphere. We don't want that, so for any shot that would put the quadrotor inside the safety sphere,

SAFETY

RESPECT A NO-FLY SPHERE

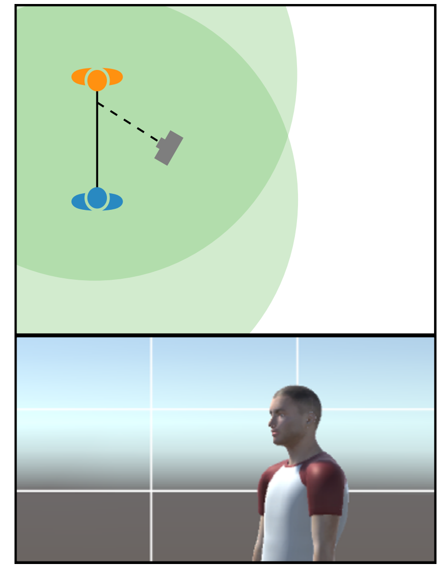

We'll approach this problem by defining a safety sphere around every person, indicated in green. We calibrated this safety sphere during flight to a value our users felt comfortable with.

Currently quadrotor cameras are mostly equipped with fairly wide angle fixed lens cameras. So for many of our shots, the default camera placement is quite close, and inside this safety sphere. We don't want that, so for any shot that would put the quadrotor inside the safety sphere,

SAFETY
RESPECT A NO-FLY SPHERE

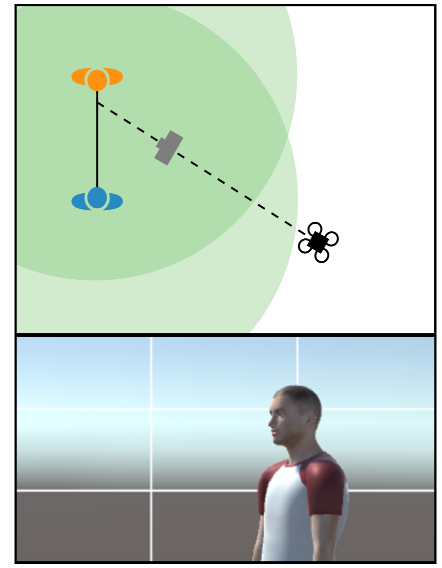

Push Quadrotor Out

We'll approach this problem by defining a safety sphere around every person, indicated in green. We calibrated this safety sphere during flight to a value our users felt comfortable with.

Currently quadrotor cameras are mostly equipped with fairly wide angle fixed lens cameras. So for many of our shots, the default camera placement is quite close, and inside this safety sphere. We don't want that, so for any shot that would put the quadrotor inside the safety sphere,

SAFETY
RESPECT A NO-FLY SPHERE

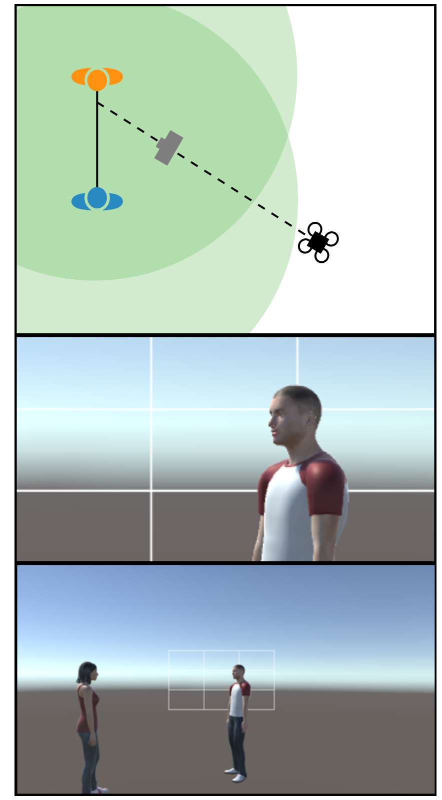

Push Quadrotor Out

Crop or Zoom Camera

We'll approach this problem by defining a safety sphere around every person, indicated in green. We calibrated this safety sphere during flight to a value our users felt comfortable with.

Currently quadrotor cameras are mostly equipped with fairly wide angle fixed lens cameras. So for many of our shots, the default camera placement is quite close, and inside this safety sphere. We don't want that, so for any shot that would put the quadrotor inside the safety sphere,

SAFETY
ADAPT CANONICAL SHOTS

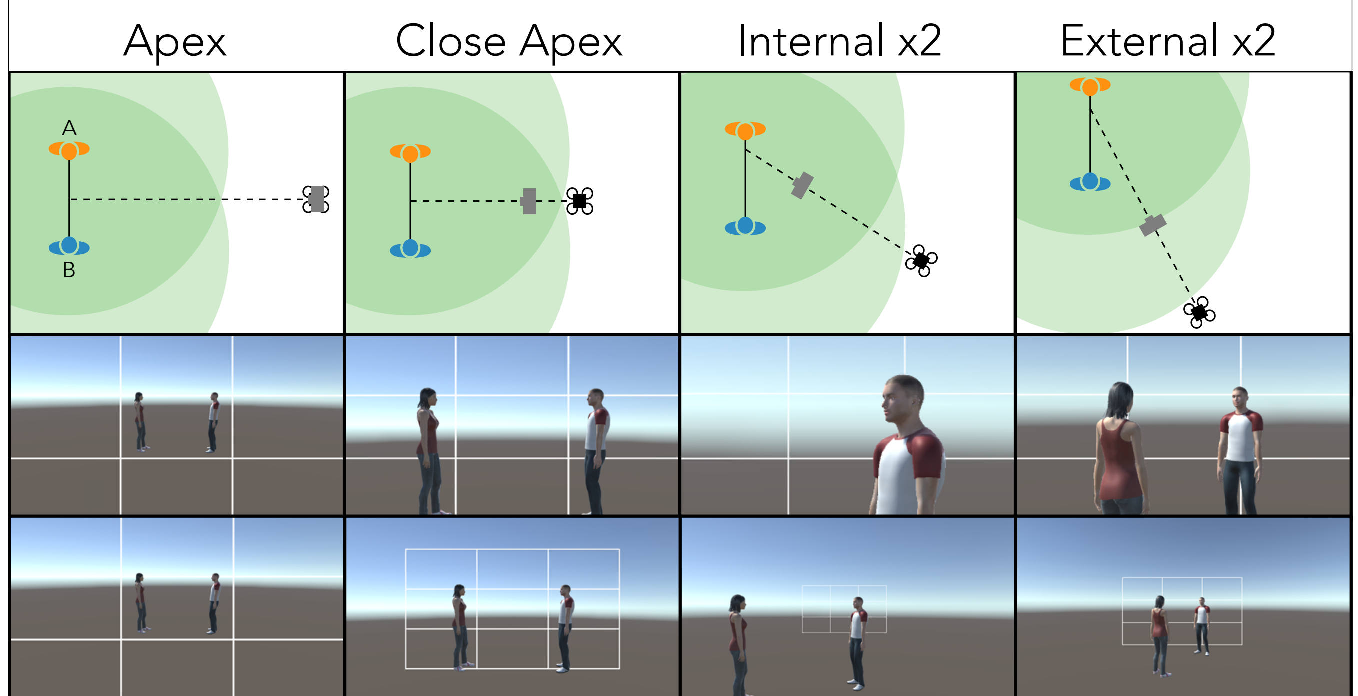

Here you can get a sense of all of our shots after we adapt them to our safety sphere. The internal shot is the closest to subjects, so gets pushed the furthers away.

Cropping does have some limitations from the resolution loss, and from our experiments it turns out we're not that great at pointing cameras yet. Since this work was done, the first consumer quads with serious zoom lenses has appeared on the market, so we think this is a pretty viable approach.

We'll run into safety again when we try to find transitions between these shots, but for the moment, I'll call these shots safe.

TRACKING

We have to know where are subjects are so we can find these camera positions. So we need some way to track them.
…

TRACKING
CAN WE USE CONVENTIONAL GPS?

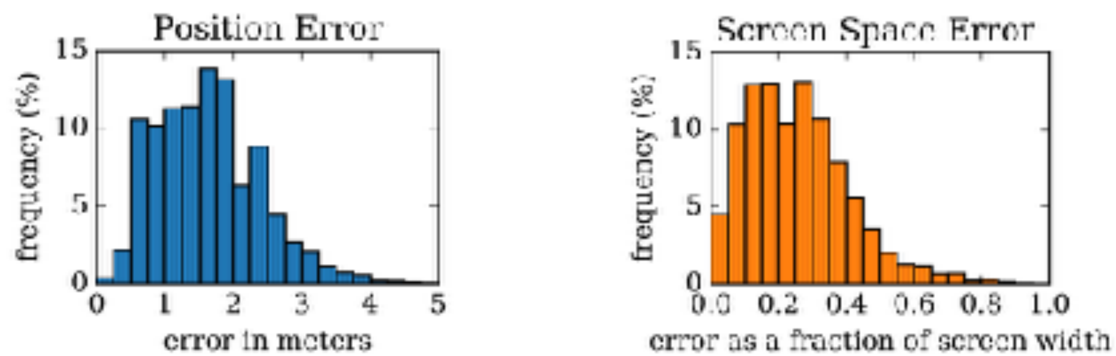

Not accurate enough!

It's common for current commercial quadrotors to have a follow-me mode, which can track a GPS like the one in your phone, or built into the remote control. We were curious whether we could localize our subjects simply by tracking the GPS, say, in their phone, or asking them to wear a GPS.

We ran a set of experiments to examine the accuracy of conventional GPS-based tracking systems for cinematography.

[CLICK] We found that today's state-of-the-art conventional GPS systems, as used on quadrotors and in cell phones, give us on average about 2m of error in position.

We then simulated capturing our canonical shots using a conventional GPS. We found that

[CLICK] on average, a person being tracked with conventional GPS is about a third of the screen width off.

That's quite significant, because we're attempting to respect the rule of thirds. If our GPS is so bad that we don't know we're in the wrong third, all our hard work to follow the rules of cinematography doesn't help much. Even worse, this means we might think we're a safe distance away from the person, when we're actually very close.

[CLICK] Concluded that conventional GPS is not accurate enough, so we looked into alternatives.

TRACKING
A CENTIMETER ACCURATE GPS TESTBED

RTK-GPS vs Conventional GPS
Tracking Position Accuracy

	Ours	Conventional
North-East CEP_{95}	0.017 m	1.68 m
Altitude Std. Dev.	0.020 m	0.108 m
Distance Error after Loop Closure	0.011 m	1.058 m

Fortunately, there's a solution. Recent advances in microprocessor technology has reduced both the cost and size of centimeter-accurate GPS systems. This GPS technology is known as Real-Time Kinematic GPS.

We ran the same experiments using this gps technology, and here's a sense of the accuracy in comparison with conventional GPS.

TRACKING
A CENTIMETER ACCURATE GPS TESTBED

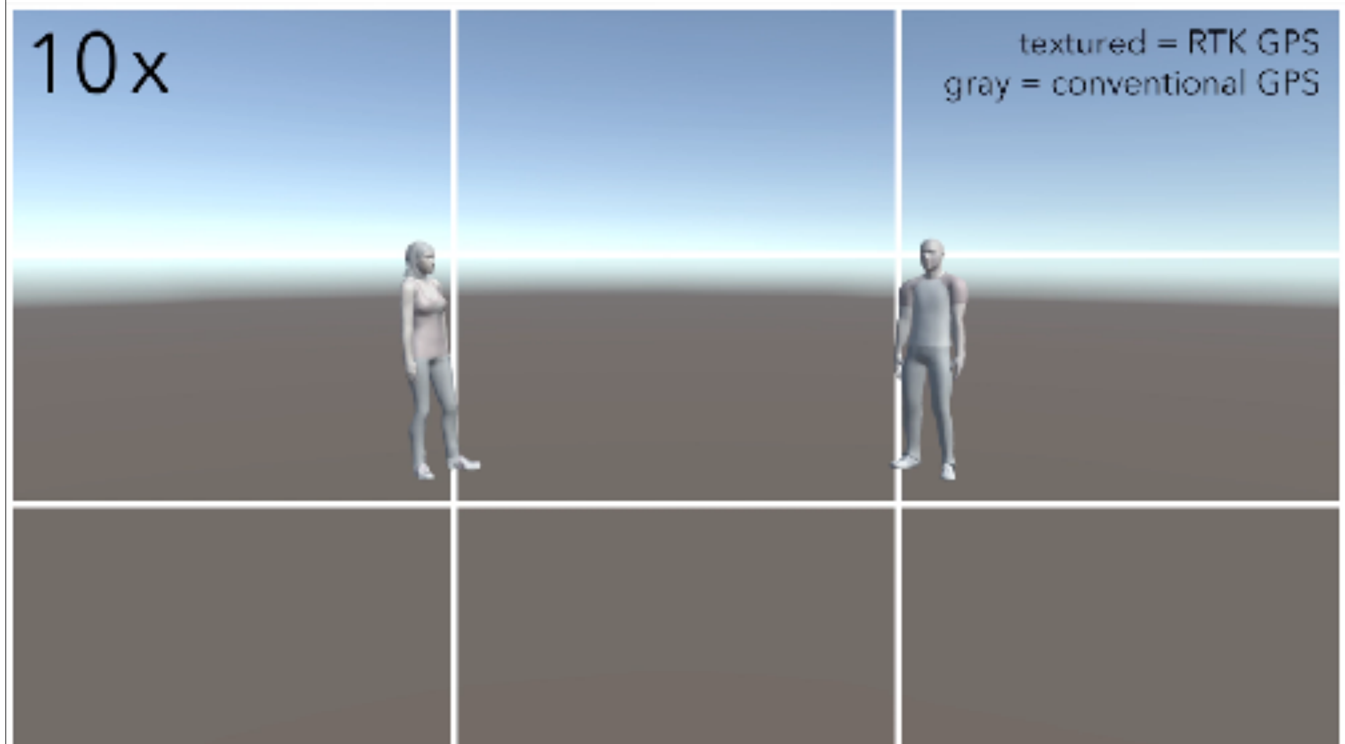

Here's a visual comparison of the two.

We fixed two conventional and two RTK gps trackers on top of two tripods. We recorded their position for a few minutes. Then we rendered this virtual preview of two subjects, as if we're tracking them. Watch how, in grey, the conventional GPS drifts all over the place, while the accurate RTK-GPS remains almost perfectly stationary over time.

TRACKING
END-TO-END SYSTEM

Using these RTK-GPS receivers, we built out an end-to-end tracking system. We also equipped our quadrotor camera with an RTK-GPS. This allowed us to place our quadrotor within about 30 centimeters of our desired position on average. In comparison, state-of-the-art indoor tracking systems that depend on having expensive motion capture rigs don’t do that much better, they get about 10 centimeter accurate quadrotor placement.

Now, when people use this testbed, we do ask our users to wear tracker helmets. In some situations this doesn’t matter much, like if you’re recording sports, where helmets are worn anyway. But it’s quite obvious in our shots.

We think this technology makes for a really great testbed. it enabled us to capture shots outdoors in real scenarios, rather than follow the path of many other researchers that limit themselves to working in a motion capture room, and it allowed us to get a sense of the accuracy we need to achieve. It’s also useful as ground truth for other methods.

We did consider a few extensions to this system.

An obvious one is using computer vision. This is a promising direction and an active area of research. There is one slightly tricky aspect of using computer vision to track multiple subjects, like we do. We have to be able to see all our subjects in our field of view! Since some of our shots are of only one person, we couldn’t just use the video feed from the main camera to track our subjects.

Since we’ve

RESULTS

SAFE STATIC SHOTS, ACCURATE TRACKING

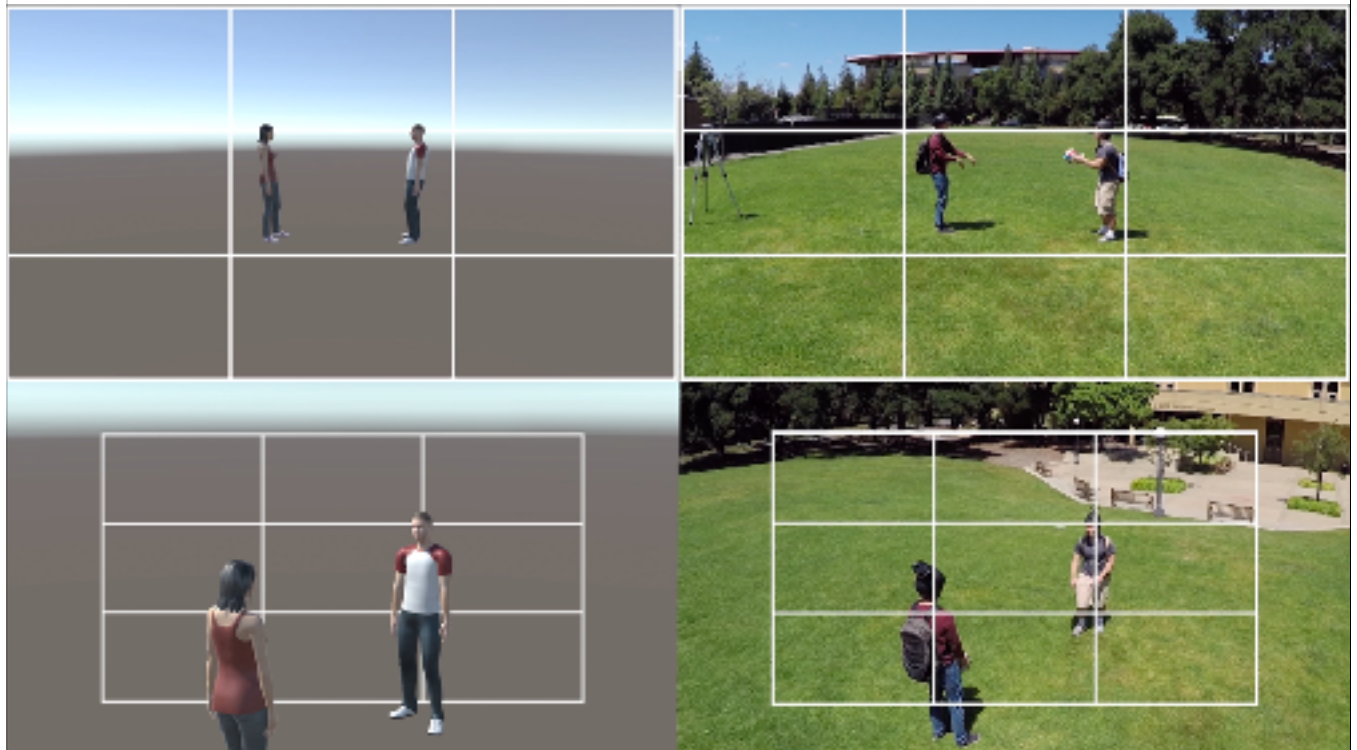

Here we're capturing two people playing catch.

On the top is an apex shot, and on the bottom is a top-down external shot.

On the left is a virtual render of the these two shots. On the right is the quadrotor faithfully capturing this composition. As you can see, our tracking and control system is accurate enough that we can actually frame people really well.

RESULTS
SAFE STATIC SHOTS, ACCURATE TRACKING

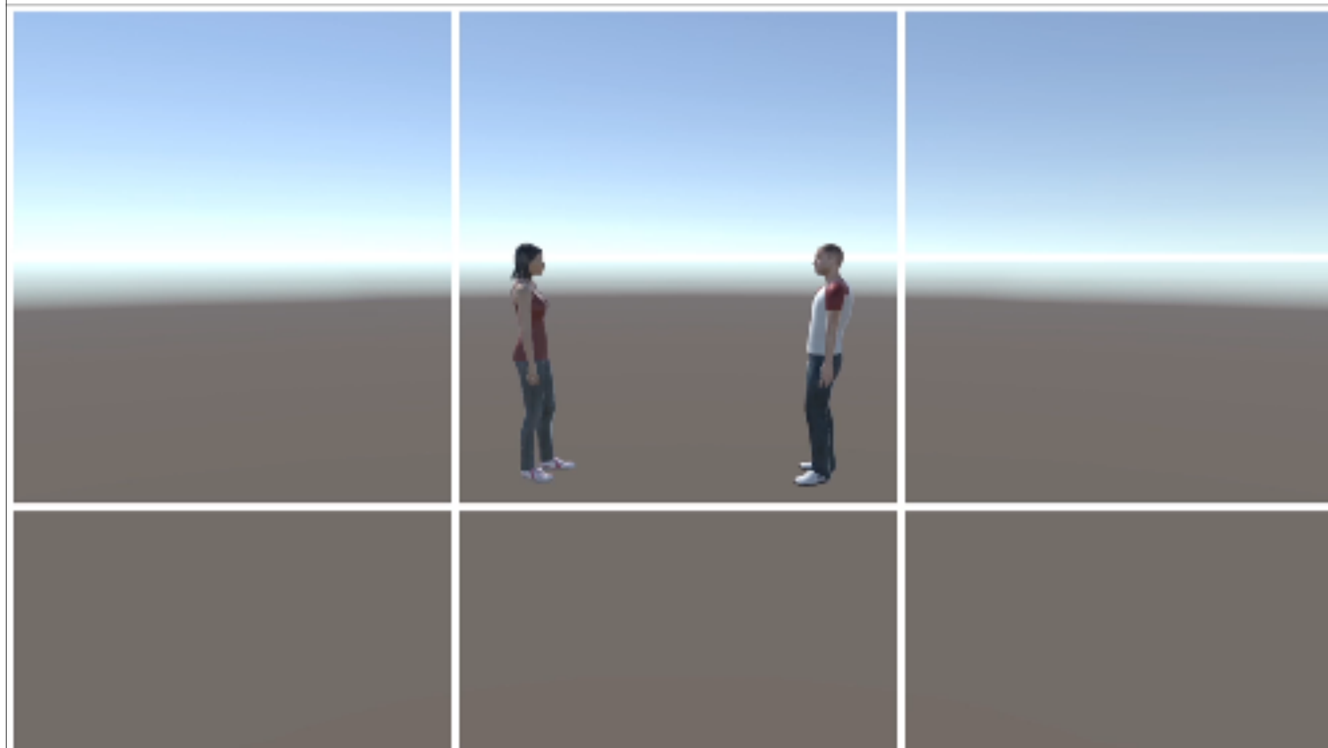

Let met show you a video of these shots.

Here’s an external shot. Notice where the two subjects fall on the Rule of Thirds line

RESULTS
SAFE STATIC SHOTS, ACCURATE TRACKING

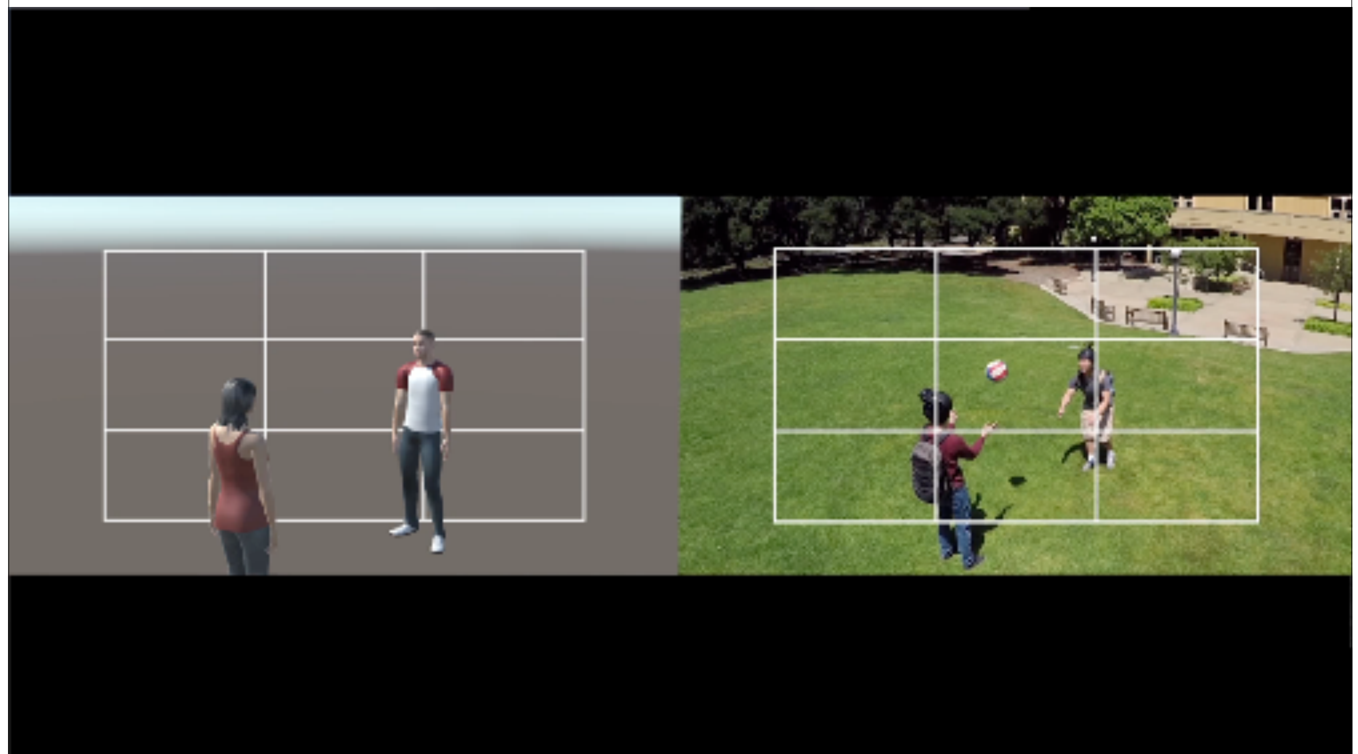

Here's the external top-down shot. You do see a little bit of drift over time, but it's really not bad for a flying machine.

RESULTS
SAFE STATIC SHOTS, ACCURATE TRACKING

Let me show you a variety of shots. It's pretty neat how much visual diversity we get out of a fairly small set of shots!

You can also see the effects of our cropping in the bottom left shot, it's slightly pixelated.

TRANSITIONS

How do we move the quadrotor camera between static shots?

Can we use the trajectories from Horus?

SIMPLY USING HORUS TRAJECTORIES DOESN'T WORK

We only guarantee respecting the physics of the quadrotor between keyframes. Previously, this was good enough since the user can add keyframes and tweak easing curves as necessary until he gets his desired shot. In this case, using only these shot templates as keyframes, means we don't control how close we get to subjects. We might fly straight through them. Similarly, just respecting the physics of the quadrotor along the trajectory doesn't mean we'll get a reasonable visual composition as we move between these predefined shot locations.

TRANSITIONS

We must move the quadrotor in a way that is

feasible

safe

produces **visually pleasing** video

and find such a trajectory **on demand, in real-time**

Three competing constraints:
- when we move between these shots, it has to be feasible, we want to keep the user safe, and it has to look good!

This is challenging since introducing obstacles means we can no longer use fast optimization methods that depend on convexity.

TRANSITIONS
USE THE TORIC SPACE METHOD?

[Lino and Christie, SIGGRAPH, 2015]

Once again we'll turn to the Computer Graphics community.

Lino and Christie introduced a new camera planning space, they call the "Toric Space". Last year at SIGGRAPH, they showed something cool and surprising. They showed how it's possible to analytically find paths with nice visual properties between viewpoints of people. Their method moves people smoothly in screen space, and tries to keep subjects in frame as much as possible, as well as respecting the line of action. So they could do it in real time no problem.

[CLICK] Here's what their results look like.

What's even better is their method trivially leads to C4 continuous paths, so it satisfies the equations of motion for our quadrotor camera!

I'd love to use their method, we can produce paths in real time between viewpoints no problem.

But they didn't respect safety!

TRANSITIONS
INSPECT THE TORIC SPACE METHOD

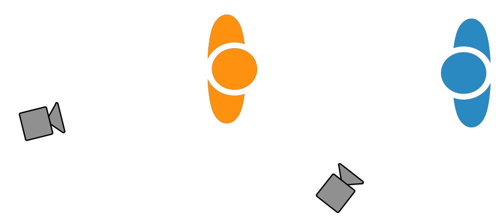

Roughly speaking, their approach is as follows:

If the want to find a camera path between these two camera positions…

TRANSITIONS

INSPECT THE TORIC SPACE METHOD

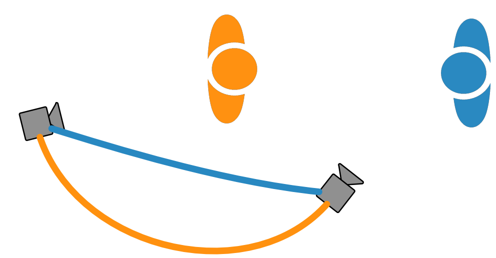

They analytically find a nice trajectory for each person, by considering only one person at a time. They do this by planning an “orbit” with respect to each person. The intuition here is, if I want to keep the person the same size and at the same point, then I spin my camera around them. I can also push in, and they’ll stay at the same point and only change size.

….

TRANSITIONS
INSPECT THE TORIC SPACE METHOD

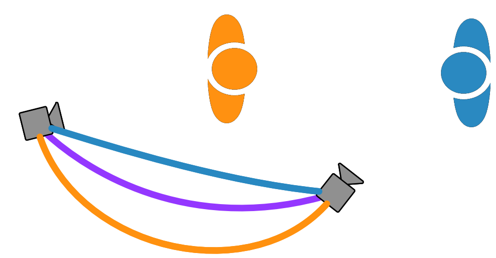

They then blend between these two trajectories in a way that nicely averages these visual effects, so you get the person moving smoothly around in your frame.

TRANSITIONS

INSPECT THE TORIC SPACE METHOD

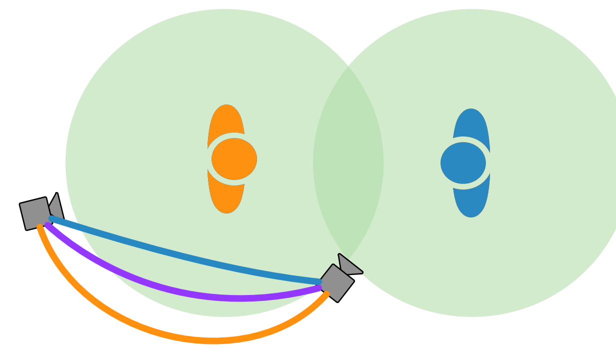

But this blend doesn't respect safety.

TRANSITIONS

INSPECT THE TORIC SPACE METHOD

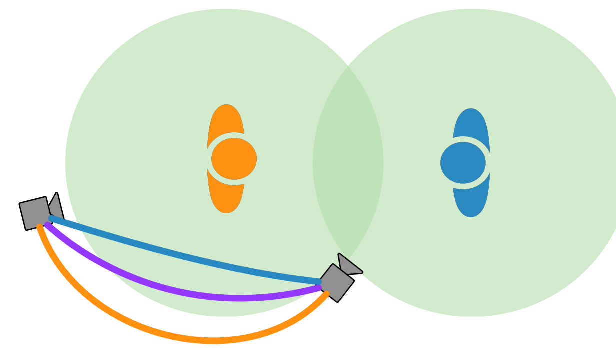

When looking at this, we made a surprising discovery. it turns out we can come up with a different blend, and we can do it in a way that respects safety and ensures feasiblity. The proof for this is a bit involved, but it turns out we can write down a small nonconvex optimization problem, and we can use it to blend these candidate trajectories. We'll make knowledge about quadrotor physics into this optimization, so we can find feasible…

TRANSITIONS
INSPECT THE TORIC SPACE METHOD

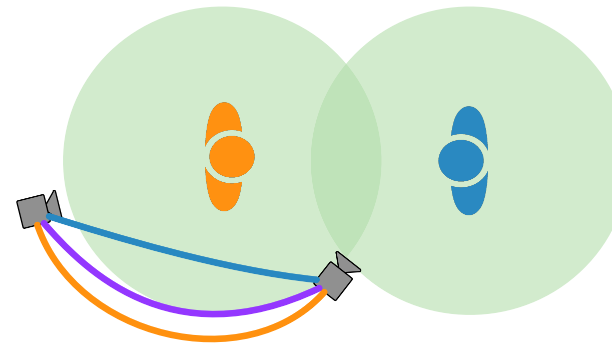

and safe paths. Our approach produces the same paths as theirs if safety isn't an issue. Our paths also produces smooth visual behavior.

The problem is quite well behaved in practice, and we can solve it in real time.

Let me show you what it looks like when we fly a camera along these trajectories.

RESULT: SAFE, FEASIBLE AND VISUALLY PLEASING TRANSITIONS

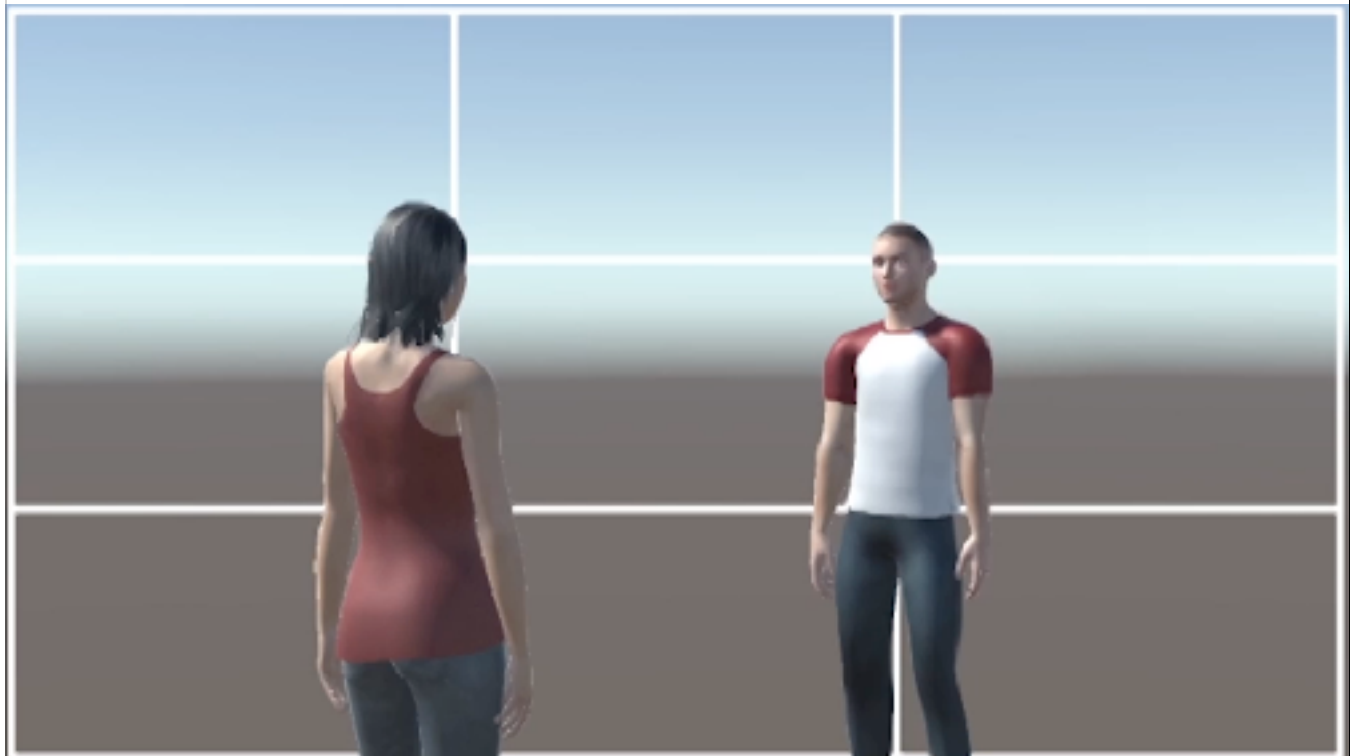

Let me show you what these end up looking like.

Here's a transition between two external shots, rendered in our virtual environment……

and here we're capturing it with a real quadrotor camera.

RESULT: SAFE, FEASIBLE AND VISUALLY PLEASING TRANSITIONS

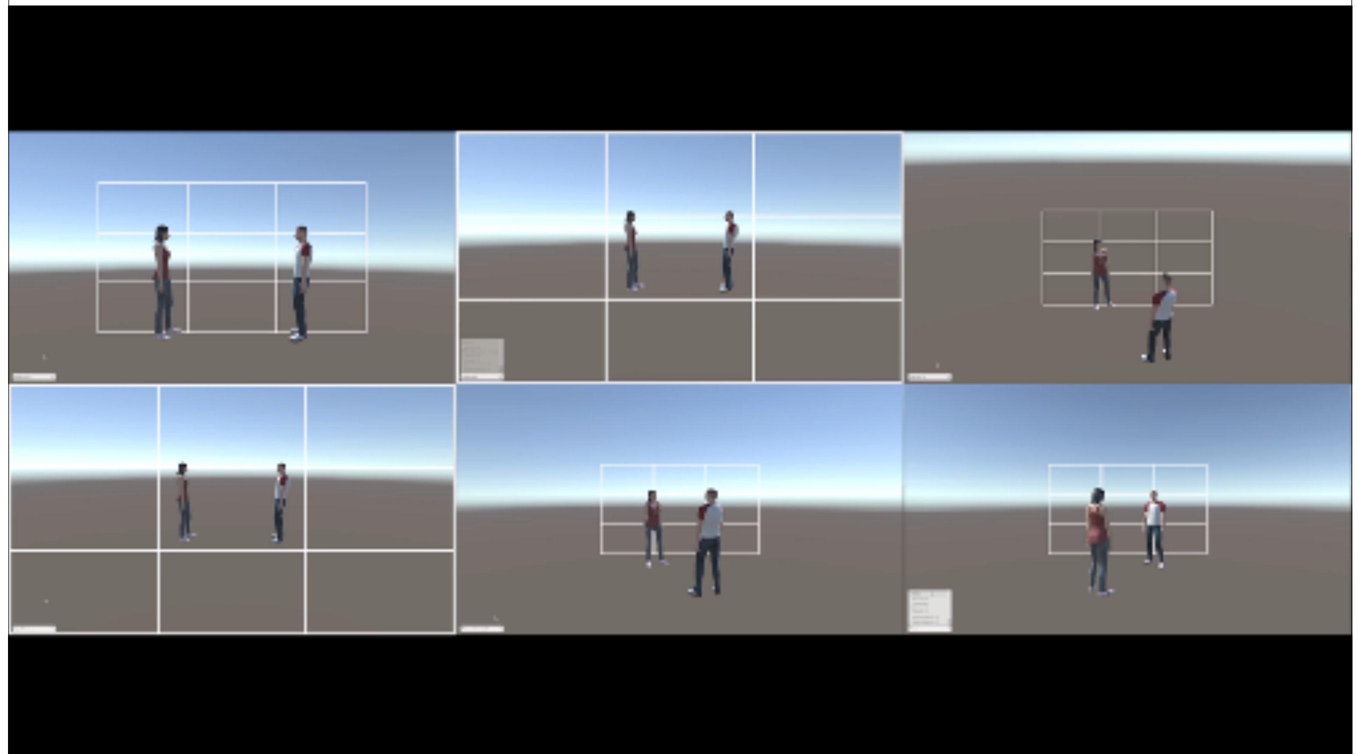

Here's a whole set of transitions. As you can see, we also smoothly animate the crop along these paths.

Now, if we put this all together, we can film something like ….

WE CAN AUTOMATICALLY CAPTURE SEQUENCES OF WELL-COMPOSED SHOTS

And if we put it all together, we can capture this simulated graduation sequence. I hope the irony is not lost on my committee.

There is no pilot, there's nothing set up ahead of time. All I'm doing is clicking which type of shot I want next. We can set a drone loose, and capture a whole range of compelling shots automatically! Totally awesome! And notice: We're using WELL-ESTABLISHED rules that produce good film, and having a quadrotor execute them!

SUMMARY

We automated **high quality visual compositions**

We adapted canonical shots to ensure **safety**

We validated, and then **tracked** our subjects, using state-of-the-art RTK-GPS

We introduced a **real-time trajectory planning algorithm** to find visually pleasing, safe and feasible transitions.

In summary,

[CLICK] we automated high quality visual compositions

[CLICK] we validated RTK-GPS for cinematography, then we accurately tracked our subjects and controlled our quadrotor using it.

[CLICK] We introduces a real-time trajectory planning algorithm to transition between shots of people, in a way that keeps them safe, is visually pleasing, is feasible, and can be computed in real time.

WITH THE DRONE CINEMATOGRAPHER, THE USER ACTS AS A DIRECTOR.

With the Drone Cinematographer, the user acts a bit like a director. The user commands well-known shots from the language of film. The Drone Cinematographer acts like a cinematographer, placing the camera and composing the shot.

OVERVIEW

Compose shots using classic 3D Animation primitives, adapted to respect quadrotor camera physics [SIGASIA 2015]

Horus

A Tool for Shot Planning

Compose shots in real time using visual composition principles from filmmaking [arXiv 2016]

Drone Cinematographer

A Tool for Filming People

Let’s recap.

Horus removed the cognitive burden of real-time control. Remained fully expressive. Made high quality cinematography easier to express than control sticks. I think of this a little bit like the Photoshop of drone cinematography. <CLICK>

Drone Cinematographer made high quality visual compositions accessible to novices, by providing high quality preset shots to the user. I think of this a little bit like the Instagram of drone cinematography. <CLICK>

CONTRIBUTIONS

Horus

[SIGASIA 2015]

Proved C4 continuous look-at look-from trajectories and progress curves satisfy Quadrotor Camera equations of motion

Showed that Computer Graphics primitives enable users to capture high quality quadrotor shots, even if they're novices

Drone Cinematographer

[arXiv 2016]

Adapt canonical cinematography shots to ensure safety

Track subjects with high enough accuracy for cinematography using RTK-GPS

Showed how to find visually pleasing and safe trajectories
for quadrotor cameras filming people

And to summarize, we made the following technical contributions

<CLICK> Horus
<CLICK>….
<CLICK> Drone Cinematographer
<CLICK>…

CINEMATOGRAPHY-FIRST INTERACTION

We **automated the role of the pilot**

Provided tools that enabled novice and experts alike
to **capture impressive cinematography**

by incorporating **quadrotor demands**
into **cinematography primitives**

We demonstrated this idea of cinematography first interaction.

[CLICK] No longer do users have to manually fly a quadrotor camera, or mess with mission planning tools.

[CLICK] We provided tools that enabled novice and experts alike to capture impressive cinematography.

[CLICK] and we did this by incorporating quadrotor demands into cinematography primitives…

THE FUTURE IS BRIGHT!

Plenty of other cinematography primitives

- “Through the lens” drone control?

Exciting new tracking and control systems

- Ultra-Wideband short range tracking? Multi-sensor fusion?

Hyper-agile quadrotors!

- 10g and 20g quadrotors? 400ft in under 1 second

Creative uses of drones beyond cinematography

Plenty of other cinematography primitives! For example, “Through the lens” control, has been developed in computer graphics. Lets you specify screen positions for objects, and it’ll attempt to move the quadrotor to capture that as closely as you can. As you expect, this also doesn’t work off-the-shelf with quadrotors. Using this we can imagine an easy interface for users to implement their own shots. This is also a possible way to support an arbitrary number of subjects in the Drone Cinematographer.

Tracking system we built is a great test-bed for many other outdoor tracking systems. For example, new ultra-wideband radio technologies promise centimeter-accurate tracking without having to rely on GPS satellites, at least over short distances. Fusing multiple sensors, like Lidar, GPS, and computer vision, could ultimately let us track people without placing anything on them.

Even more exciting, for me, is that the Drone Racing world is pushing the limits of quadrotors. We’re already very close to building quadrotors that can generate 10gs of force. What if we can do 20g’s. That gets us from sitting on the ground to 400ft, the FAA legal limit for quadrotors, in under 1 second. Imagine the types of shots we can get with that.

IMPACT

Our work is already having an impact in the commercial world.

All our code is open-sourced and freely available.

Since then…

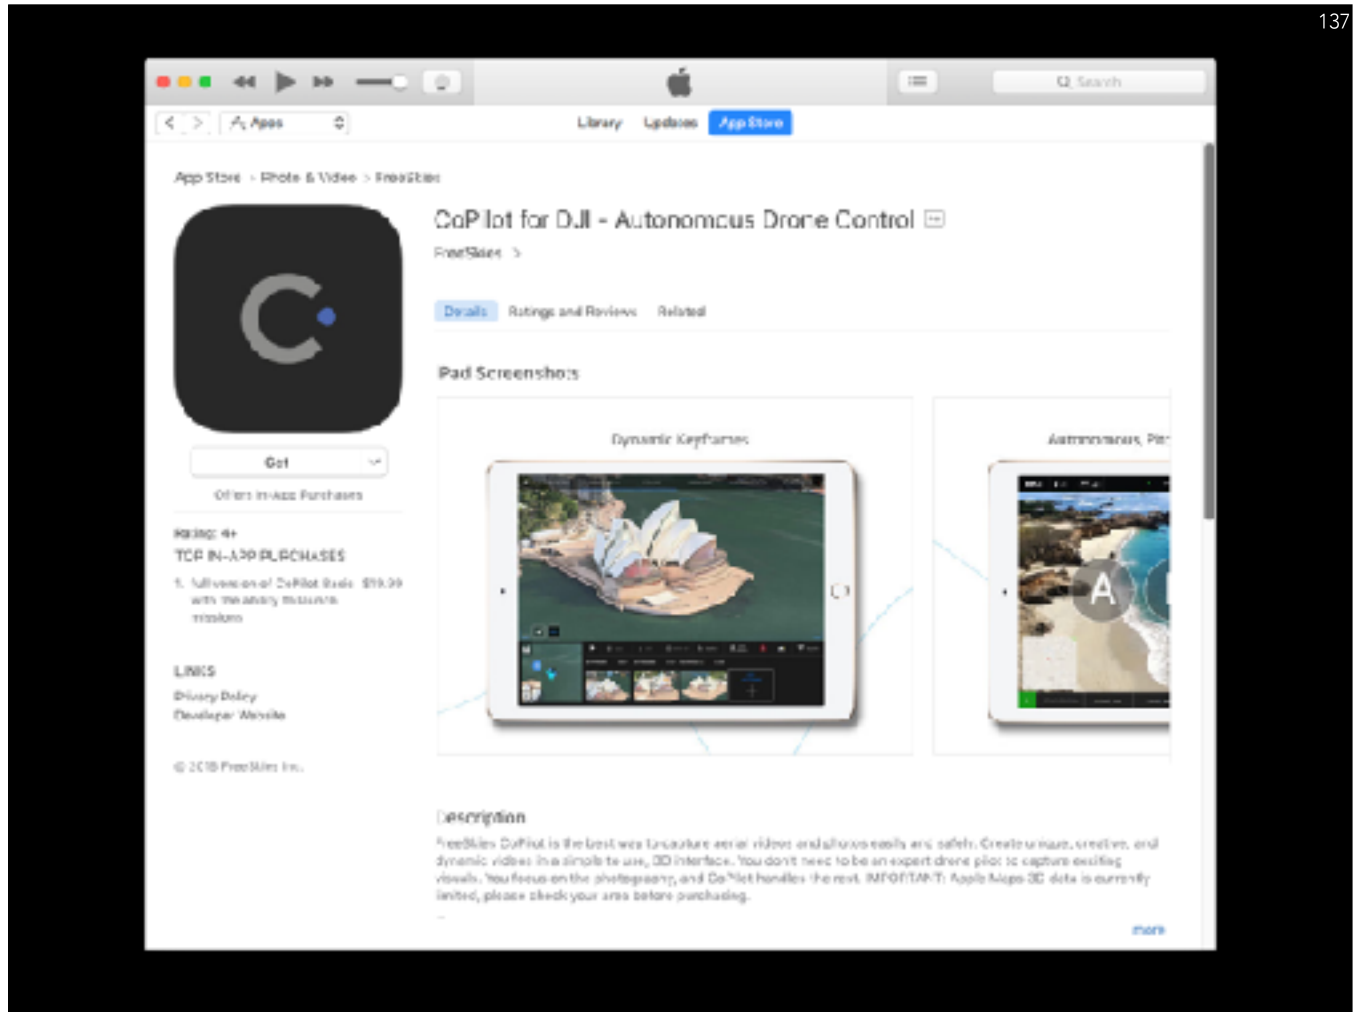

Our friend Andy and his team at Freeskies have recreated Horus as an iPad app.

5000 downloads, 1000+ shots captured in the last 90 days.

You can all be trying this out before leaving the room today.


3D
location
launch

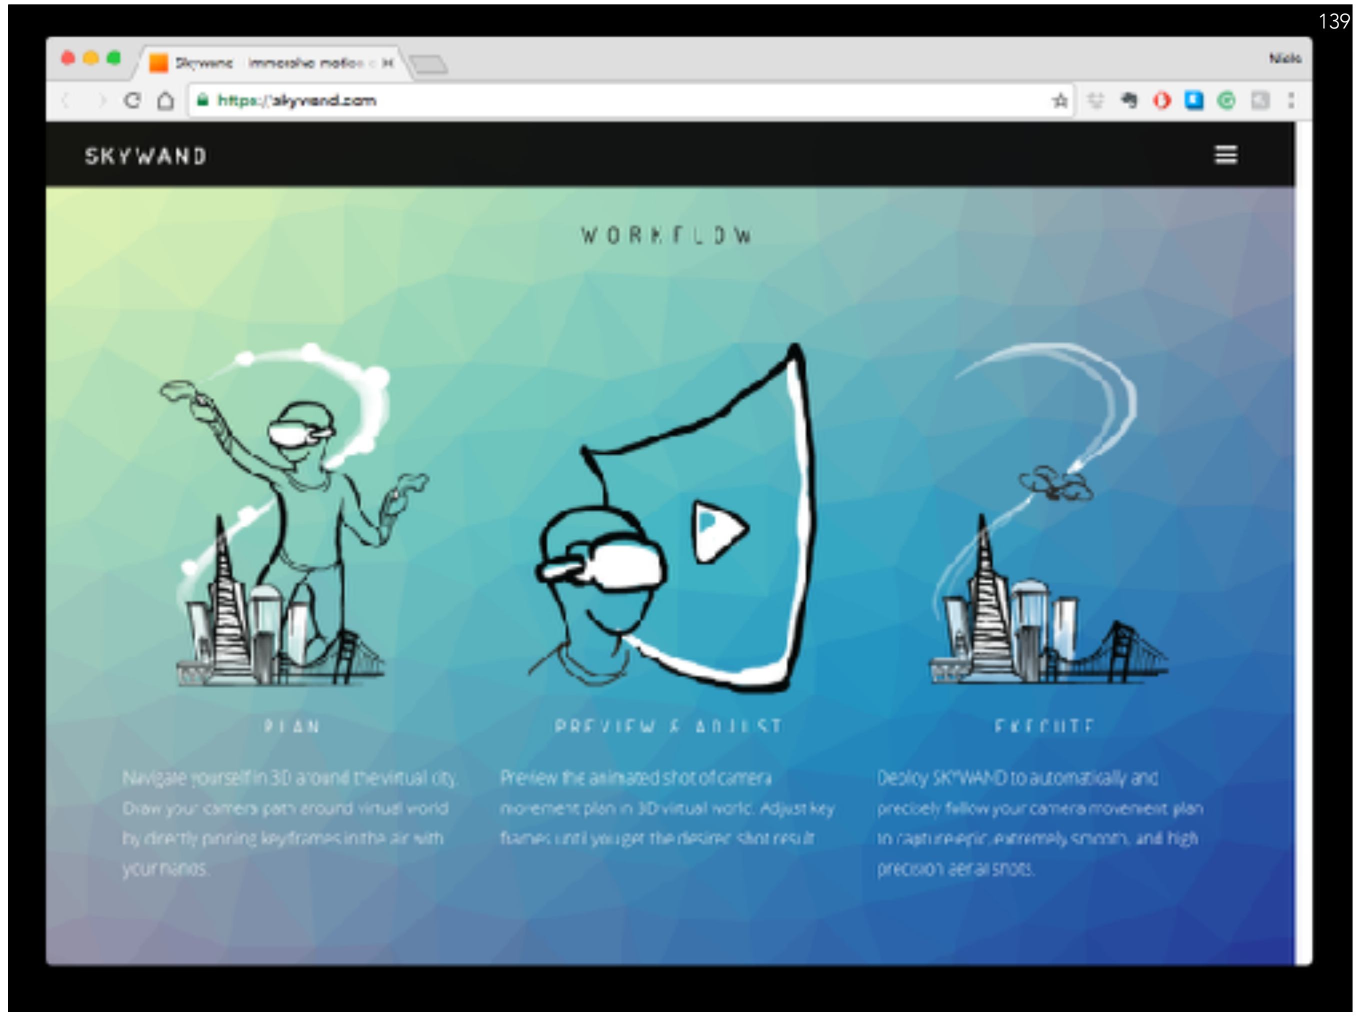

Stanford graduate Amber started Amber's Garage, a VR startup. His first software product is Skywand, a VR implementation of Horus. He also extended our codebase to work with DJI quads.

Here’s how he sets up keyframes in his virtual environment. The paths are planned using our codebase…

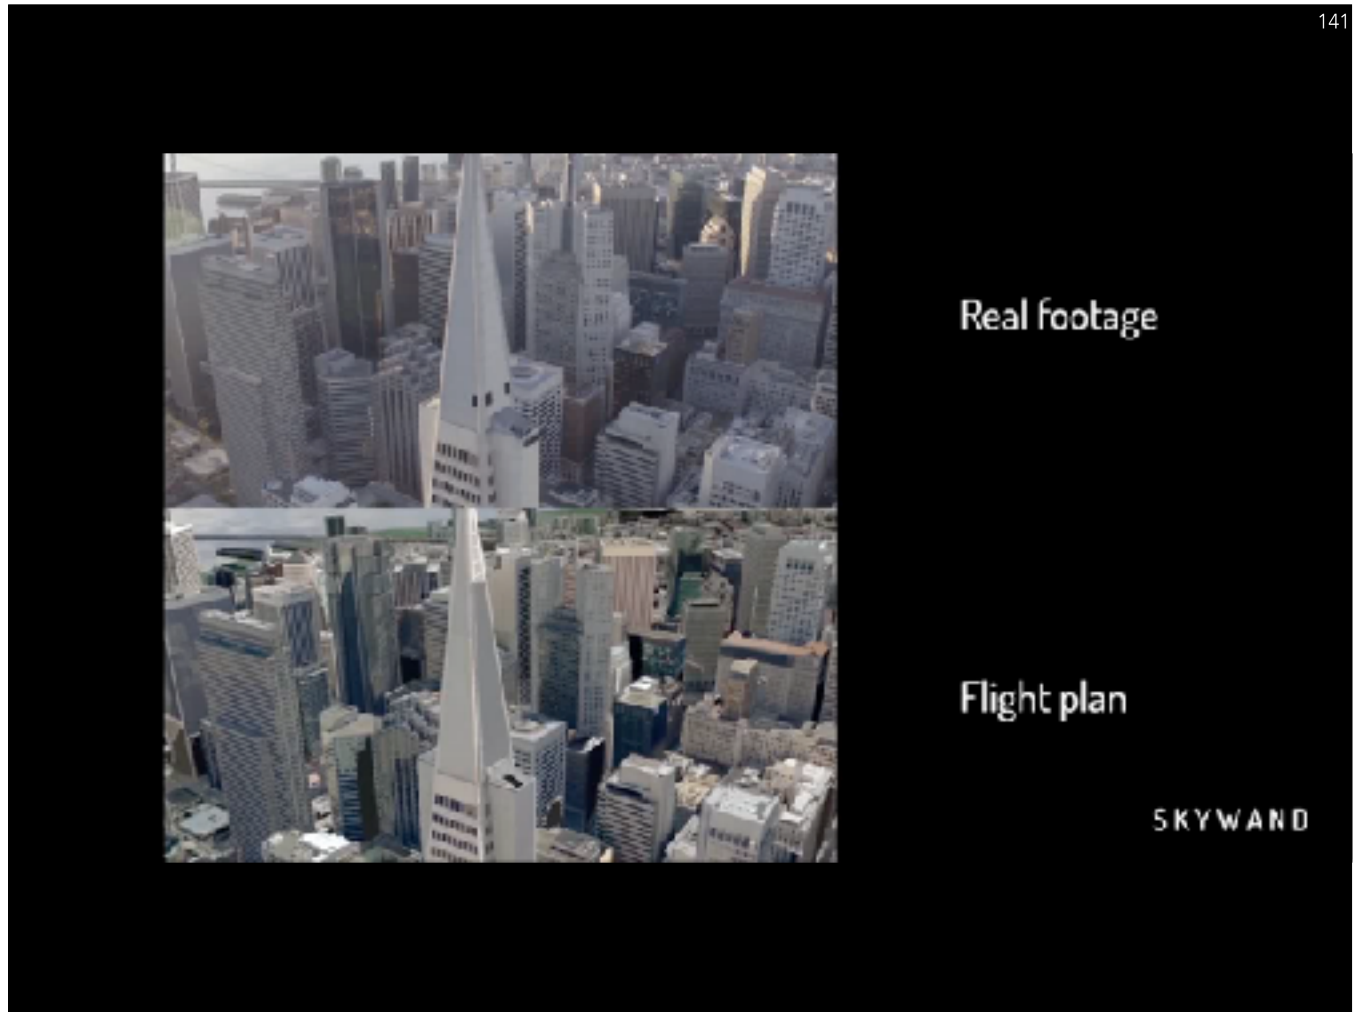

and here’s some of the epic shots he's capturing using today’s latest and greatest quadrotors.

There’s also interest from multiple other industry partners that we’re busy talking to.

That’s it for the technical part of my presentation

FEELINGS

Now it’s time to talk about feelings. I have a lot of people to acknowledge and thank for this work.

THANK YOU

Pat
Committee - Stu, Dave, Juan, Maneesh, Mac
Collaborators - Mike R, Jane E, Anh T, Dan G, James L, Zach D, Juan A, Eric D, Andrew K, Sebatian B, Alex A, ….
Industry Partners - All the 3DR folks, Swift Nav folks, Adobe folks, DJI
User Study Participants
Stanford UAV Club & Drone Cinematography community - Trent, Eli, Eric C, George K,…
Students - Anh, Harrison, Vicky, Eli, Kat, Stephanie, Noa, Jorge
Fellow Inmates & Colleagues in the Graphics lab, Stanford, Pixar, Yahoo, …
Burning Man families the Dusty Connection, the Shady Waffle, Monks of Funk
Friends off the Farm Gleb, Ashley, Marcello, Yizhuo, Patrick, Trisha, Eric, Trent, Fergus, Carrie, Sher, Jennifer, Hendrik in SA, …
EECS House @ Cal
Mentors - Henry Chamberlain, Annelie Starke, Lona Antoniadis, so many more
Silicon Graphics branch, South Africa - Neil McGowen, …
Extended Family - Ackermans, Du Toits, Oosthuizens, Jouberts, …
Girlfriend Liz, who lives life so well that Death will tremble to take her

The only reason I'm standing here today is because of the incredibly people that's carried me. There's far too many to mention by name but I'm going to try anyway and were i fail I'm gonna try to thank entire groups.

[Click] Pat, my advisor. I've been fortunate enough to work on three very different projects over the course of my PhD, and he jumped into this crazy quadrotor project before our research community was paying any attention to drones. That kind of freedom and trust is rare. Thank you!
[Click] Committee. Especially Stu, I spent many hours over the last month hanging out at his house, and to get that kind of input from a luminary in the field is outstanding. Dave, who has been excited about this work for years and helped us get so much hardware. Juan, who also helped us get Stanford's UAV club off the ground, Maneesh, who followed me here from Berkeley, and Mac, who agreed at the last minute to come chair my defense.
[Click]Collaborators, especially Mike, Jane and Anh, without whom this project never would have happened

Last on this list, because she gets her own line, my lovely girlfriend,

THANK YOU

Mental Health

I want to separately take a second to thank those who kept me sane over the course of my PhD.
47% of graduate students struggle with depression during their PhD, 20% struggle with substance abuse, 10% seriously contemplate suicide. And we don't talk about this much, yet we see it in our very department. I know I wouldn't be here if it wasn't for those who helped me hang on to threads of sanity, especially when those threads wear thin.
So to my meditation teach, therapist, and of course all the friends, colleagues, family, thank you.

=======

*Graduate student happiness and well-being report, UC Berkeley, 2015
**Arora et al, 2016

FLORAINE BERTHOUZOZ
1984-2015

NELLIS OOSTHUIZEN
1929-2016

I also need to thank those who saw this work getting started, but will never have the chance to see it come to fruition.

And lastly…

The Joubert Clan. My brothers, Pierre Henri, Dieter and my Mom and Dad.

Look Mom, I made it, I’m all grown up.

A VISION OF THE FUTURE

"The Universe is made of Stories, not of Atoms"
- Muriel Rukeyser

Robotic systems that intelligently balance
Human Interfaces, Aesthetic and Technical Knowledge

will enable capturing your stories, while remaining immersed in your experience.

THANK YOU

I would like to close with talking about the vision that's driving this work, and I'll do that with a quote I love: "The Universe is made of Stories, not Atoms"

Currently we spend a lot of time caught up in trying to capture the moment. So much that we often forget to actually live it, and to be present in the moment.

It's my hope that, In the future, robotic systems that intelligently balance XXX…

And with this thesis, we're taking a couple of steps in the direction of that future. With that, I'll take your questions